



AWS Organizational Change Acceleration (OCA) 6-Point Framework – 3.
Envision the Future

AWS Prescriptive Guidance



AWS Prescriptive Guidance: AWS Organizational Change Acceleration (OCA) 6-Point Framework – 3. Envision the Future

Copyright © 2026 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

Table of Contents

Introduction	1
Intended audience	3
Targeted business outcomes	3
About the OCA 6-Point Framework guides	3
3.1 Change strategy and plan	5
Overview	5
Best practices	7
FAQ	7
Additional steps	7
3.2 Communications strategy and plan	9
Overview	9
Best practices	10
Getting started	11
FAQ	12
Additional steps	13
3.3 Engagement strategy and plan	14
Overview	14
Best practices	14
Engagement plan components	15
Categorizing stakeholders	18
Prioritizing and mapping stakeholders	19
Measuring success	21
Examples	21
FAQ	22
Additional steps	23
3.4 Training strategy and plan	25
Overview	25
Best practices	26
FAQ	27
Additional steps	28
3.5 Risk mitigation strategy and plan	29
Overview	29
Best practices	29
FAQ	32

Additional steps	33
3.6 Sponsor roadmap	34
Overview	34
Best practices	34
FAQ	38
Additional steps	39
3.7 Sustainability plan	40
Overview	40
Best practices	41
FAQ	41
Additional steps	42
Resources	46
References	46
Partners	46
Contributors	48
Document history	49
Glossary	50
#	50
A	51
B	54
C	56
D	59
E	63
F	65
G	67
H	68
I	70
L	72
M	73
O	77
P	80
Q	83
R	83
S	86
T	90
U	91

V 92

W 92

Z 93

AWS Organizational Change Acceleration (OCA) 6-Point Framework – 3. Envision the Future

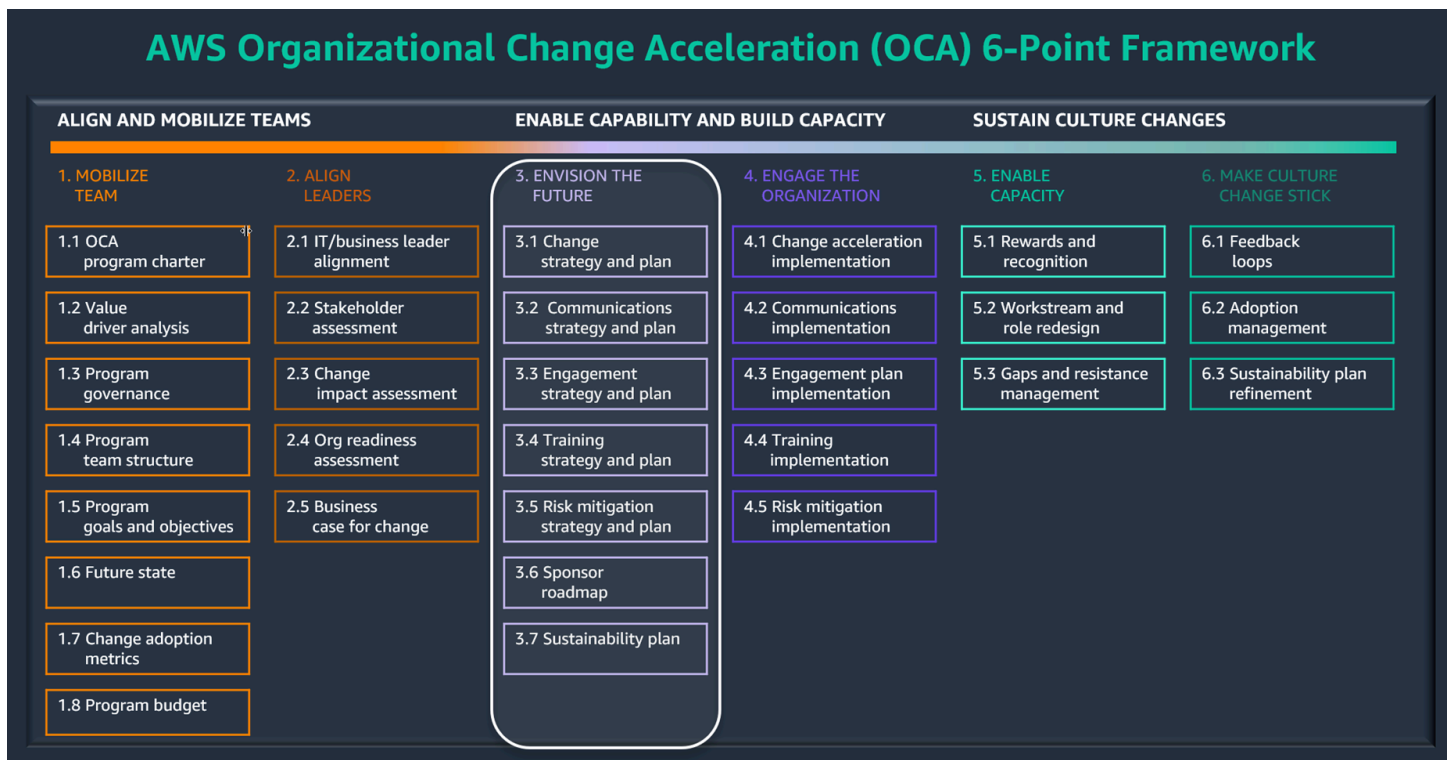
Amazon Web Services ([contributors](#))

January 2025 ([document history](#))

The AWS Organizational Change Acceleration (OCA) 6-Point Framework is intended to cover the full scope of people-related issues and challenges throughout the lifecycle of a cloud transformation, which might include migration, modernization, generative AI scaling, and innovation. This framework guides customer adoption of AWS technologies, processes, and new ways of working by:

- Identifying, aligning, and mobilizing key leaders
- Assessing and mitigating the organizational impacts of cloud transformation
- Designing change acceleration, communications, and training plans
- Developing leadership, sponsorship, and culture strategies

The framework's six points align with an agile sprint cadence, from program initiation through sustainable long-term change. The following diagram shows these six points and their subpoints.



The third point, *Envision the Future*, helps create a change acceleration strategy and plan to communicate, train, and engage the employees of the organization in their cloud adoption journey. It contains seven subpoints:

- [3.1 Change strategy and plan](#). Summarize the strategy and plan for full change acceleration across workstreams. Ensure alignment and understanding across workstream leads and leadership.
- [3.2 Communications strategy and plan](#). Promote awareness, understanding, and desire for the cloud future state.
- [3.3 Engagement strategy and plan](#). Engage key stakeholders to enable the organization to move to a defined cloud future state..
- [3.4 Training strategy and plan](#). Make sure that identified stakeholders obtain the necessary knowledge, skills, and abilities to implement future cloud processes.
- [3.5 Risk mitigation strategy and plan](#). Proactively identify and eliminate or control people-related risks associated with the cloud future state.
- [3.6 Sponsor roadmap](#). Make sure that sponsors and leaders are aligned on cloud objectives and are held accountable for taking action to reduce risk and accelerate cloud adoption.

- [3.7 Sustainability plan](#). Support the desired future-state behaviors and organizational structures beyond the initial phase of cloud adoption.

This guide discusses each subpoint of *Envision the Future* in detail.

Intended audience

This guide targets leaders who are responsible for accelerating cloud transformation. Following these recommendations will help minimize risks and maximize value.

Targeted business outcomes

The *Envision the Future* phase of the AWS OCA 6-Point Framework contributes to the following outcomes:

- **Organizational alignment:** Envisioning the future establishes an ongoing partnership between organizational structures, business operations, processes, workforce, and culture. This enables rapid adaptation to market conditions and the ability to capitalize on new opportunities.
- **Cloud acceleration:** Envisioning the future identifies and minimizes impacts to people, culture, roles, and organizational structure when moving from the current state to a future state. This accelerates the adoption of new ways of working.
- **Cloud fluency:** Envisioning the future builds digital acumen to leverage the cloud effectively and to accelerate business outcomes.
- **Transformational leadership:** Envisioning the future mobilizes leaders to drive transformational change and to enable outcome-focused, cross-functional decision-making.

About the OCA 6-Point Framework guides

This guide is part of a set of publications that cover the OCA 6-Point Framework, which is a programmatic and evidence-based organizational change adoption framework.

The content set includes a comprehensive set of templates, guidelines, supporting artifacts, assessments, accelerators, and tools that are designed to accelerate cloud transformation. We recommend that you start with the [overview](#) to understand the framework and its six points, and then consult the following individual guides for detailed discussions of each point.

1. [Mobilize Team](#)
2. [Align Leaders](#)
3. Envision the Future (this guide)
4. [Engage the Organization](#)
5. [Enable Capacity](#)
6. [Make Culture Change Stick](#)

For a comprehensive set of cloud transformation strategies, guidance, and resources, see [Accelerating cloud transformation](#).

3.1 Change strategy and plan

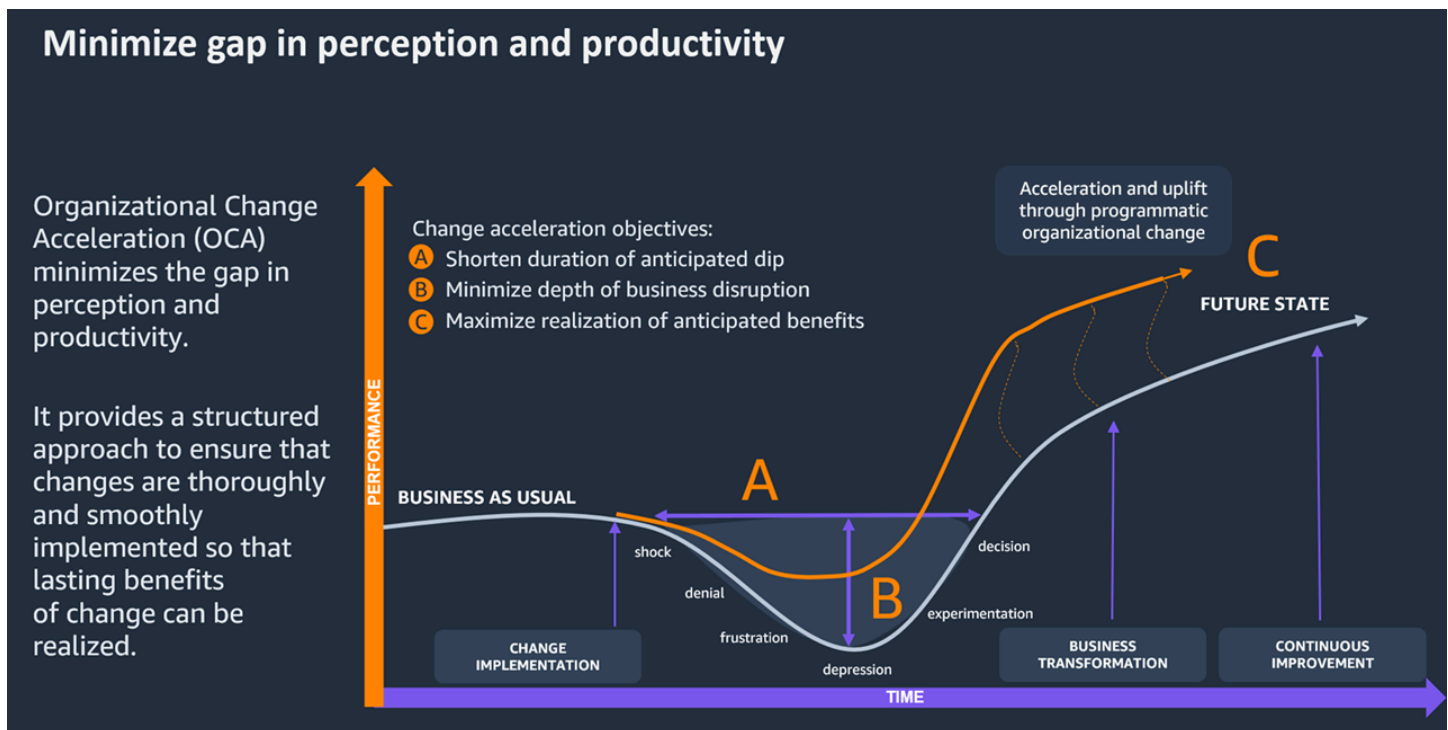
Overview

The change acceleration strategy and plan provide a structured approach to delivering the right change tactics to the right people at the right time during cloud transformation. They ensure that changes introduced by the cloud are accepted with minimal disruption and maximum results. The strategy outlines how the organization will address changes in operations, technology, structure, and processes, and inform decision-making throughout the transformation.

According to Accenture's [Modern Cloud Champions](#) study, organizations that manage the people side of change effectively are six times more likely to meet or exceed project objectives. A well-designed change strategy can deliver significant benefits. It can:

- Minimize risk, performance dips, and business disruptions.
- Ensure business continuity and maintain customer service levels.
- Secure leadership alignment and commitment.
- Prepare all impacted audiences for change.
- Foster culture transformation and new ways of working.
- Increase stakeholder awareness, engagement, and understanding.
- Position the organization for ongoing success and adaptability.

New systems and strategies can be highly disruptive to an organization. A well-formulated strategy brings the project or change to life, and describes who it will impact in the organization and how. It minimizes gaps in perception and productivity, as shown in the following diagram.



[Accenture's study](#) revealed that emphasizing people as much as technology in cloud transformation leads to significant improvements. In their study, this approach led to:

- 2.2x improved organizational agility and innovation
- 2.2x stronger business and IT collaboration
- 1.9x faster cloud migration
- 1.7x improved customer experience
- 1.7x achieved cost savings
- 1.7x increased speed or efficiency to launch new lines of business
- 1.6x increased cloud adoption
- 1.4x better use of data for decision-making

You should develop the change acceleration strategy at the program's outset, and create, review, and update the accompanying plan at key milestones, phases, releases, or epics throughout the program. OCA requires high integration among partners. To ensure the strategy's success, you must establish and maintain effective partnerships among human resources (HR), the cloud transformation team, executive sponsors, leadership, external vendors, and other relevant partners.

Best practices

Key best practices for a change acceleration strategy include the following:

- Align the change strategy with the cloud transformation's strategic vision and business case to add emphasis and increase consistency.
- Interview key leaders and stakeholders early in the planning process to gain insights and build buy-in.
- Assess stakeholder alignment periodically throughout the program to ensure continued support.
- Integrate various partners to ensure consistent objectives, timelines, and motivations across the organization.
- Partner with enabling functions such as HR, training, finance, and cross-functional leaders to leverage their expertise and resources.
- Develop a flexible strategy that can adapt to emerging challenges and opportunities.
- Incorporate metrics to measure the effectiveness of change initiatives.

FAQ

Q. Who should be involved in this activity?

A. This activity must be conducted with the executive sponsor, project leader, change leader, internal change team liaison, internal communications, and HR.

Q. What are the inputs to the change strategy and plan?

A. Key inputs include the strategic vision, [business case for change](#), discovery documentation, OCA project charter, [stakeholder assessment](#), leadership interviews, internal change leadership (if available), and internal communications (if available).

Q. When should the change acceleration strategy and plan be created?

A. Initiate the design of the strategy at the onset of the program. Create, review, and update the accompanying plan at key milestones, phases, releases, or epics within the program.

Additional steps

To begin building the change acceleration strategy and plan, follow these steps:

1. Review the strategic vision and [business case](#) to ensure alignment.
2. Analyze discovery assessment findings and outputs.
3. Review the [organization readiness assessment](#) to identify gaps and opportunities.
4. Interview executive sponsors, migration team leadership, internal communications, and HR.
5. Consider aligning the strategy with your existing internal change methodologies (if applicable).
6. Leverage OCA goals and objectives to guide strategy development.
7. Incorporate defined future state findings into the strategy.
8. Establish change acceleration governance and project team structures.
9. Draft a comprehensive change acceleration strategy.
10. Review and validate the strategy with the leadership team to ensure buy-in.
11. Obtain sign-off on the final change acceleration strategy.

By focusing on these key elements and best practices, you can develop a robust change strategy for your organization that accelerates cloud adoption, maximizes business value, and ensures a smooth transition for all stakeholders who are involved in the transformation journey.

3.2 Communications strategy and plan

Overview

As your digital transformation begins to affect your organization and is no longer contained at a leadership level, the stakes for communications become high. You must solidify your messaging on the case for change and develop milestone plans to define the audiences and appropriate channels to deliver and receive communications. The upfront design of a communications strategy will define who the audiences are, what the messages need to convey, where target audience members will receive those messages, when impacts and calls to action will be revealed, and how communications will be monitored and measured for reach and effectiveness.

Early in the planning phases, developing a communications strategy is important, and its implementation could be the determining factor in how well the transition is received and behaviors are adopted by your audience. For example, are team members aware and prepared for a data center exit? Do managers know when they can plan to support their employees' upskilling plans? Is it clear to leaders what happens if cloud adoption stops or stalls?

The purpose of the communications strategy is to provide a thoughtful, structured approach to delivering the right messages to the right people at the right time throughout the course of your cloud transformation. There are subtle differences between a communications strategy and a communications plan. Here is how these documents are defined at Amazon Web Services (AWS):

- *Communications strategy* – A document that expresses the goals and methods of an organization's outreach and communication activities.
- *Communications plan* – Detailed information about communication activities that address the strategy and achieve the objectives listed in the strategy. The plan describes each activity with information such as delivery date, intended audience, detailed message, media type, creator, approver, and messenger.

An effective communications strategy and plan can deliver significant benefits. They can:

- Increase awareness and understanding of the cloud transformation across the organization.
- Align stakeholders on the vision, goals, and progress of the transformation.
- Reduce resistance to change by addressing concerns and highlighting benefits.
- Accelerate the adoption of new processes and technologies.

- Improve employee engagement and motivation throughout the transformation.
- Enhance trust and transparency between leadership and employees.
- Support a smooth transition to the new ways of working.

A series of documents and assessments will help to inform the foundation of the communications strategy and plan. Here are some key inputs and outputs.

Inputs	Outputs
Strategic vision	Communication guiding principles
Business case	Message identification
Stakeholder assessment	Stakeholder prioritization matrix
Communication methods assessment	Vehicle and media analysis
Change impacts	Communication matrix
Cloud transformation milestones and status updates	Communication activities (work plan and approach)

Best practices

- Start early. Develop the communications strategy and plan in the early planning phases of the cloud transformation.
- Understand and follow the communication process to best architect essential messaging and communication activities.
- Orchestrate a series of interviews with key stakeholders to identify the level of communications required for the plan.
- Align and time communications with key project milestones and decision points.
- Use a variety of communication channels to reach different audiences effectively.
- Customize messages for different stakeholder groups based on their needs and interests.
- Encourage two-way communications. Create opportunities for feedback and dialogue throughout the transformation process.

- Regularly assess the effectiveness of communications and adjust the strategy as needed.

Getting started

To create a communications strategy, start with a communication needs assessment with direct input from the cloud transformation strategic vision, the [business case](#), and the [stakeholder assessment](#).

The communications development process consists of the following steps:

1. Assess communication needs.
2. Develop communications strategy and plan.
3. Develop communications.
4. Validate communication content.
5. Distribute communications.
6. Collect feedback.
7. Measure effectiveness.

A communications strategy contains the following components.

Component	Description
Communication objectives	The rationale and the importance of communicating effectively throughout the digital transformation or migration.
Communication guiding principles	Core values to observe in communications. For example, be direct, honest, and open; demonstrate integrity.
Communication vehicles	The channels to use for communications. For example, websites, blog posts, videos, vlogs, social media messages, emails, online newsletters.

Component	Description
Key communications messaging	Project status announcements, major milestone accomplishments, initial message to explain the reason for change.
Target audiences and stakeholders	End-users, suppliers, company leaders, IT managers and supervisors, business leaders, general audiences, steering committee, cloud migration team, transformation management office, and so on.
Communications approach and work plan	A visual representation of all communication activities that are built into the plan.
Organizational roles and responsibilities	The communication-related tasks and duties of executive sponsors, steering committee members, process owners, champions, project team, internal communications team, and HR team.

Note

Depending on your organization, you might add more components to the communications strategy.

FAQ

Q. When should you use the communications strategy and plan?

A. Developing a communications strategy and plan is important, and how well you implement these could be the determining factor in how well your organization receives the transition and adopts target behaviors. Communications strategy and planning efforts start at the beginning of your cloud transformation project—typically as soon as project goals and preliminary milestones are established. During this phase, gaps, change impacts, and teams and employees who are

affected by the migration are identified. When you build your communications strategy, follow the communication process within your organization to best architect cloud messaging and activities.

Q. Who is involved?

A. Executive sponsors, digital transformation leader, internal communications, and HR teams are typically involved in creating the communications strategy and plan.

Q. What are the inputs to the communications strategy and plan?

A. Inputs include the strategic vision, business case, stakeholder assessment, communication methods assessment, change impacts, and cloud transformation milestones and status updates.

Q. What are the outputs of the communications strategy and plan?

A. Outputs include communication guiding principles, message identification, stakeholder prioritization matrix, method and media analysis, communication matrix, and communication activities (work plan and approach).

Additional steps

To begin to create the communications strategy and plan, complete these tasks:

1. Collect information from discovery documents, including the business case and stakeholder assessment.
2. Conduct interviews with the executive project sponsor and project leadership team.
3. Conduct interviews with internal communications teams.
4. Conduct interviews with HR to understand potential impacts to future state roles.
5. Conduct interviews with functional process area leaders.
6. Conduct interviews with outward-facing groups such as sales.
7. Evaluate all documentation and information, and follow up with key stakeholders as needed.
8. Build your communications strategy deck.

By focusing on these key elements and best practices, you can develop a comprehensive communications strategy and plan that support the cloud transformation journey, engage stakeholders effectively, and drive the adoption of new ways of working.

3.3 Engagement strategy and plan

Overview

The engagement strategy and plan outline a systematic approach that describes specific ways in which individuals, stakeholder groups, or organizations will address the changes caused by the cloud transformation. The primary intent of the engagement plan is to keep all key stakeholders committed to, and focused on, the desired business results of the cloud transformation. Identifying stakeholders and engaging them appropriately throughout the change process are critical to the success of the project.

The engagement strategy and plan heighten the involvement within and outside the cloud transformation team. They ensure that the right people receive the right information, so they can participate at the right time and in the right way. They work as a forcing function to proactively manage the pace and amount of change that each stakeholder group must undergo to avoid overload.

An effective engagement strategy and plan can deliver significant benefits. They can:

- Increase stakeholder buy-in and commitment to the cloud transformation.
- Identify and mitigate potential roadblocks early in the process.
- Enhance the organizational capability for change.
- Maximize the potential for a successful transition to cloud adoption.
- Improve alignment between different stakeholder groups.
- Accelerate decision-making processes.
- Foster a culture of collaboration and shared responsibility.

Best practices

An engagement strategy and plan actively involve stakeholders and can help identify, manage, and avoid potential roadblocks. These documents result in additional organizational buy-in, commitment, and capability for cloud transformation, and further maximize the potential for successful cloud adoption.

The OCA team goal for this activity is to:

- Determine where stakeholders stand and create an engagement plan to influence them in a way that aligns with the cloud migration vision.
- Secure strong leadership alignment and support.
- Collaborate with HR and the internal change team, if available, to understand the organization's change practices used in the past.

Engagement plan components

The following illustration shows the key components of the engagement strategy and plan, what each component does, and the target audience.



The following table provides more information about each component.

Component	Activities
Leadership alignment and sponsorship	<ul style="list-style-type: none"> • Share and reinforce the IT vision with business teams. • Establish an operating rhythm of communication and dialogue.

Component	Activities
	<ul style="list-style-type: none"> • Provide intranet support assignments (for example, blogging). • Plan events (for example, lining up business speakers). • Provide regular updates at leadership meetings. • Communicate and celebrate successes. • Identify new improvement initiatives. • Gain alignment on priorities.
IT education and development	<ul style="list-style-type: none"> • Provide rich and collaborative content on the intranet, including: <ul style="list-style-type: none"> • Leader blogs, discussion forums, articles, external sites, professional associations, news, case studies • Communities of practice • Produce educational materials, including: <ul style="list-style-type: none"> • Quick, animated e-learning with voice-over • Capability-related educational concepts, case studies, application activities <p>IT leadership should receive the materials and coaching kit 48 hours before training.</p> • Design formal development programs for employees, starting with VPs.

Component	Activities
Specialized communications	<ul style="list-style-type: none"> • Launch the IT vision through: <ul style="list-style-type: none"> • Web meetings with the IT community • Video introduction by the chief technology officer (CTO) on the intranet • Comments solicited in IT and business discussion forums • Provide general updates on progress through intranet announcements, emails, town halls, conference calls, and web meetings. • Share success stories and quick wins. • Host a road show to promote communications, learning, and community building.
Improvement initiatives	<ul style="list-style-type: none"> • Establish an organization-wide time for learning, for a specific duration and on a scheduled cadence. • Launch a net new rewards and recognition program that's specific to the enterprise transformation.

Each component also includes ongoing change acceleration monitoring that involves the following activities:

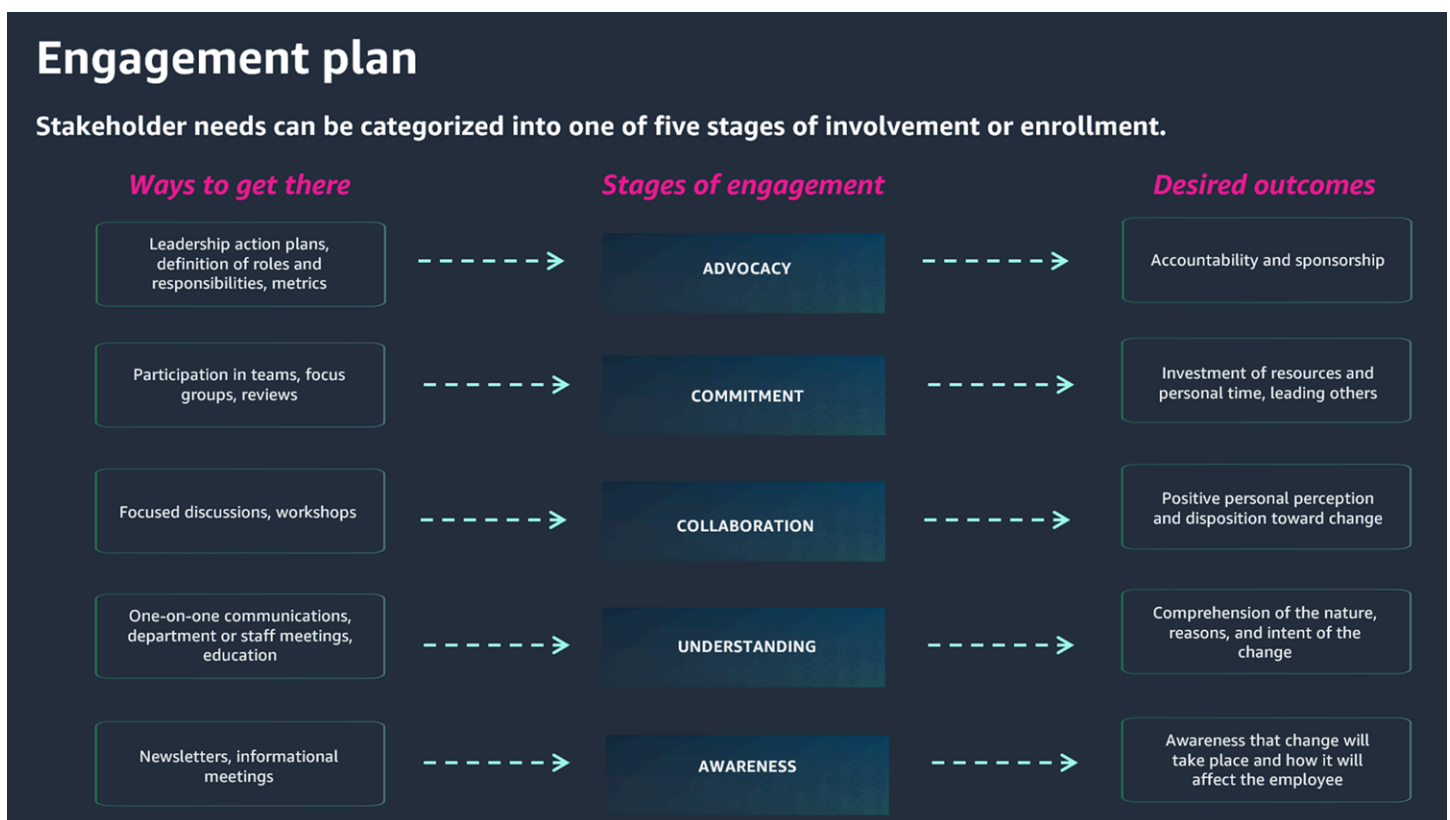
- Monitor and measure change awareness, understanding, and acceptance.
- Measure overall program progress and effectiveness.
- Develop, implement, and refine change and initiative plans.
- Identify new initiatives to enable change.

Categorizing stakeholders

After you develop the plan, place each stakeholder into one of the five stages of involvement or engagement (from least involved to most involved):

- **Awareness:** Stakeholders are aware and understand the purpose and progress of change.
- **Understanding:** Stakeholders have a sound understanding of the benefits and implications of change.
- **Collaboration:** Stakeholders support the change, believe it is worthwhile, and would act if prompted.
- **Commitment:** Stakeholders proactively communicate and take actions required to support the change.
- **Advocacy:** Stakeholders own initiatives and work to improve and sustain performance.

The following illustration describes ways to achieve these stages and desired outcomes.




Each stage requires a unique communication objective and mechanism to effectively engage the organization, as shown in the following illustration.



Prioritizing and mapping stakeholders

Stakeholder prioritization and mapping should occur after you conduct a [stakeholder assessment](#). The OCA team needs to build and maintain strong relationships with these stakeholders. The team can use the following matrix and place stakeholders in the appropriate quadrant based on how critical they are to transformation success and the degree of impact to the stakeholder. After this mapping, the OCA team can develop a strategy to build and maintain the relationships.



The quadrants are:

- **Monitor and respond.** Stakeholders in this quadrant are neither highly influential nor greatly affected by the changes, but they hold a stake in the results. These stakeholders require minimal communication activities; mass communications are usually sufficient. The main objective is to monitor their feedback to avoid problems.
- **Keep informed.** Stakeholders in this quadrant are affected significantly by the outputs of the transformation but have less influence over others and less potential to disrupt the process. Communications with these stakeholders should be strongly proactive and preemptive, and more influential stakeholders (from other quadrants) should influence their acceptance.
- **Maintain confidence.** Stakeholders in this quadrant have significant influence over others and therefore have the potential to disrupt the process. For this reason, it's important to anticipate their objectives and adverse reactions when you plan communications. The project's impact on this group is lower, so there is less need to involve them in development. Communications should aim to sustain and expand their support, but need not be as resource-intensive or frequent as the next quadrant, because the need for them to change their behavior is lower.

- **Win them over.** Stakeholders in this quadrant are highly influential, and the outputs of the transformation work have significant impact on their future work processes and behaviors. The possibility of this group to disrupt the project is potentially very high. Involve them in communications, keep or develop them as allies, and emphasize frequent personal contact and face-to-face communications.


Measuring success

Measuring the success of engagement activities, plans, and the strategy is critical to the success of the overall cloud transformation. Here are some mechanisms and measurements that you can use to rate the success of the plan:

- Change readiness surveys (champions and end-users)
- Change acceleration scorecard
- Preparation and readiness sessions (evaluations)
- Training assessments
- Attendance reporting and curriculum completion
- Implementation (go or no-go) decisions

Examples

The following illustrations provide examples of change activity plans that could be derived from the engagement strategy and plan.



Change activity plan – example 2

ID #	Change activity	Description	Tools provided	Approximate timing	Frequency
1. ●	Change touchpoints	<ul style="list-style-type: none"> Discuss good practices for information sharing, what's working, and what could be improved; note change risks. Assist change team that is developing solutions to change risks. 	Planned touchpoints	Beginning on date x/y	Biweekly or as needed
2. ▲	Staff or shift meeting updates	<ul style="list-style-type: none"> Provide update on current project activities and upcoming events. Answer questions of end-user constituents. 	Project update content	ASAP after kickoff	Biweekly or as needed
3. ?	Hosted Q&A sessions and office hours	<ul style="list-style-type: none"> Host Q&A sessions on location for leaders, managers, and employees to answer questions about the IT cloud and to share information about the program. Make certain hours of the week available for people to contact you with questions about the program or their responsibilities. 	FAQ	October	Monthly or as needed
4. ✉	Awareness and engagement email communications	<ul style="list-style-type: none"> Provide informal communications to the network after large project milestones to drive awareness. For example: <ul style="list-style-type: none"> Drive traffic to the Ask IT cloud site for information, FAQ Drive traffic to the IT network news articles when published 	Templates	Milestone-driven	As needed
5. ○	Lunch and learn discussions	<ul style="list-style-type: none"> Host information session to provide an overview of what is changing. 	Overview material	November	1-2 times before go-live
6. N/A	Ad-hoc Q&A	<ul style="list-style-type: none"> Be available to answer questions as they come up. Funnel questions to change team or IT OpEx team if answers are unknown. 	FAQ	As needed	As needed

FAQ

Q. Why are the engagement strategy and plan valuable?

A. These deliverables heighten involvement within and outside the cloud transformation team, ensure that the right people receive the right information at the right time, and proactively manage the pace and amount of change for each stakeholder group. By actively involving stakeholders, they help identify, manage, and avoid potential roadblocks and increase organizational commitment and capability for cloud transformation .

Q. When do you use them?

A. Use an engagement strategy and plan after you complete the preliminary work of assessing stakeholders, creating the [change strategy and plan](#), and developing the [communications strategy and plan](#). These documents can drive ongoing support and leverage the influence of stakeholders.

Q. Who should be involved in this activity?

A. Participants should include the executive sponsor, cloud leader, OCA leader, HR lead, chief architect, data lead, security lead, operations lead, training lead, finance lead, infrastructure leaders, lines of business leads, and the internal communications team.

Q. What are the inputs to this strategy and plan?

A. Inputs include the strategic vision, business case, Migration Readiness Assessment (MRA) outputs, people acceleration project charter, stakeholder assessment (analysis), interviews with the executive sponsor, HR, and migration leadership, and input from internal change leadership (if available) and internal communications team (if available).

Q. What are the outputs of this activity?

A. The outputs of this activity are an approved engagement strategy and plan.

Additional steps

To create the engagement strategy and plan, follow these steps:

1. Review the strategic vision and business case.
2. Review the findings and outputs from discovery.
3. Review the stakeholder assessment.
4. Review and expand the stakeholder list.
5. Identify the change impact (for example, on roles and responsibilities, or training) by role.

6. Group roles into categories based on change impact.
7. Map each stakeholder group to the prioritization matrix by defining existing and desired states.
8. Define what each stakeholder group needs to know.
9. Determine the timing for providing the information to each stakeholder group.
10. Define the key messages for each stakeholder group, based on the commitment model:
 - Goals
 - Measurable objectives
 - Additional information to back up key messages
11. Review and expand on the communication methods list.
12. Determine the communication methods that are preferred for each stakeholder group.
13. Identify the messengers and message for each stakeholder group.
14. Determine how best to measure communication effectiveness and how often this assessment should occur.
15. Develop the work plan by using the communications calendar as the baseline.
16. Implement the communications plan and measure the results.
17. Manage the feedback process, and refine or redesign the approach as required.
18. Draft the engagement strategy and plan.
19. Review and validate the engagement strategy and plan with the leadership team.
20. Sign off on the engagement strategy and plan.

By focusing on these key elements and best practices, you can develop a comprehensive engagement strategy and plan that support your organization's cloud transformation journey, ensure stakeholder commitment, and maximize the potential for successful cloud adoption.

3.4 Training strategy and plan

Overview

Training is crucial for preparing your cloud migration, modernization, or transformation team to understand and perform their jobs in a cloud-transformed future state. It provides the formal structure, instruction, and practice that help users learn new processes and technologies. A well-designed training strategy ensures that employees can operate confidently in the new cloud environment.

The training strategy and plan defines:

- Target audiences
- Training methods
- Content
- Timelines
- Facilitators
- Logistics

AWS takes a data-driven approach to training requirements. You can identify your organization's cloud skills gaps by using the free [AWS Learning Needs Analysis \(LNA\)](#) tool, which helps you create targeted, cost-effective training plans.

Training for new cloud skills is often a blend of different formats, including on-demand, virtual instructor-led, in-person instructor-led, hands-on labs, [game days](#), and [immersion days](#). An effective training strategy delivers significant benefits:

- Accelerates adoption of new cloud technologies and processes
- Reduces errors and improves operational efficiency
- Enhances employee confidence and job satisfaction
- Improves organizational agility and innovation capabilities
- Increases return on investment (ROI) from cloud technologies
- Mitigates risks associated with skill gaps

- Supports talent retention by providing growth opportunities

Best practices

- Align with business objectives. Make sure that the training strategy supports overall cloud transformation goals.
- Prioritize critical competencies. Focus on the most important skills for successful cloud adoption.
- Use diverse learning formats. Blend different training methods, including:
 - On-demand courses
 - Virtual, instructor-led sessions
 - In-person workshops
 - Hands-on labs
 - Game days
 - Immersion days
- Take advantage of data-driven insights. Use tools such as the [AWS LNA](#) to identify skill gaps and to tailor training plans.
- Develop a timeline. Create a training sequence that aligns with the cloud transformation roadmap.
- Evaluate knowledge acquisition. Implement assessment methods such as certifications or practical demonstrations.
- Continuously update your plan. Regularly review and update the training plan to keep pace with evolving cloud technologies and organizational needs.

Key components of a training strategy and plan are:

- Training objectives: Define expected outcomes for participation, skill growth, certifications, and specific capabilities.
- Training guiding principles: Establish decision-making boundaries that address funding, requirements, and sourcing.
- Training methods: Determine delivery methods such as in-person, virtual, computer-based, self-paced, or a combination.
- Target audiences: Identify key roles, departments, and users for training.

- **Training activities:** Create a timeline that's aligned with the cloud program. Include specific user stories in project tracking tools.
- **Organizational roles and accountabilities:** Define responsibilities for implementing, monitoring, and measuring training and communicating results.

FAQ

Q. Why is the training strategy valuable?

A. A training strategy aligns knowledge, skills, and capabilities with organizational goals while demonstrating the value of training investments.

The [AWS LNA](#) produces learning recommendations for continued growth. It enables you to use the data you collect on the workforce, training activities, and knowledge growth in a practical way.

If you don't have a well-articulated training strategy, employees who aren't familiar with the requirements for operating effectively in the cloud and how it differs from operating on premises have to rely on their own resources to gain that knowledge. This can result in unproductive time spent on training and knowledge acquisition that aren't related to cloud operations. A clear training strategy and plan provide an individualized roadmap for all impacted employees that covers the skills they will need to successfully operate in a cloud environment.

Q. When would you schedule this activity?

A. Begin training strategy and planning efforts at the start of your cloud transformation project, after you determine gaps and change impacts, and identify affected teams and employees.

Q. What is the AWS LNA?

A. The AWS LNA is a free self-assessment tool that identifies an organization's cloud skills gaps. Employees complete an adaptive survey, and AWS experts use the results to create targeted, cost-effective training and certification plans.

Q. Who should be involved in this activity?

A. Participants should include the executive sponsor, cloud leader, OCA leader, HR lead, chief architect, data lead, security lead, operations lead, training lead, finance lead, infrastructure leaders, and lines of business leads.

Q. What are the key inputs to this activity?

A. Inputs include the [stakeholder assessment](#), training assessment, discovery documentation, and [change impact assessment](#).

Q. What are the key outputs of this activity?

A. Outputs include training guiding principles, audience analysis, training plan, training roles and responsibilities, training objectives, and training budget.

Additional steps

To create the training strategy and plan:

1. Interview internal training teams to understand specific needs and appropriate delivery methods.
2. Consult HR to understand potential impacts on future state roles.
3. Interview functional process area leaders to understand the specifics for their area.
4. Consult outward-facing groups to understand potential training impacts on customers and suppliers.
5. Evaluate all documentation and information, following up with key stakeholders as needed.
6. Build your training strategy document by following this structure:
 - Training objectives
 - Training guiding principles
 - Training methods
 - Target audiences
 - Training activities (work plan and approach)
 - Organizational roles and accountabilities

By focusing on these elements and best practices, you can develop a comprehensive training strategy for your organization that supports cloud transformation, addresses skill gaps, and enables employees to thrive in the new cloud environment.

3.5 Risk mitigation strategy and plan

Overview

People-related issues can become risks or blockers that impede the start or scaling of the cloud journey. A risk mitigation strategy and plan offer a structured approach to provide visibility into these issues, dismantle roadblocks, and accelerate change. Common people-related issues include:

- Misalignment between leaders about cloud objectives
- Prioritization differences related to timelines and resource allocation
- Communication breakdowns between siloed functions or management layers
- Cloud skill shortages throughout the workforce

Mitigating these risks saves companies effort, time, and money, and reduces organizational friction. These risks can take a toll on employees if they aren't properly managed.

An effective risk mitigation strategy delivers significant benefits:

- Accelerates cloud adoption by proactively addressing potential roadblocks
- Improves project timelines and budget adherence
- Enhances stakeholder alignment and communications
- Reduces employee stress and turnover related to transformation challenges
- Increases the overall success rate of cloud initiatives
- Provides a structured approach to continuous improvement


Best practices

- Review the cloud strategy and plan for desired outcomes and timelines.
- Align with the project manager on overall issues and the risk mitigation process.
- Develop an ongoing risk identification process.
- Establish dimensions for risk categorization, such as vision and clarity, culture, commitment, communications, retention and engagement, and skills and capability.

- Assess the severity of risk and probability of occurrence.
- Develop a risk tracking and evaluation tool (see the example table later in this section).
- Document people-related issues that might pose a risk to timely completion of people transformation deliverables.
- Look across the program to see how technical, budgetary, and timing risks will impact people and create people-related risks.
- Handle sensitive or confidential risks appropriately, and communicate these only to a small circle of people who need to know.
- Track the mitigation and closure of people-related risks over the course of the cloud program to evaluate their impact on achieving desired cloud outcomes. For example, an impact statement might be: "15 high-severity risks were identified and mitigated; if these risks had not been mitigated, the cloud journey would have been delayed by approximately 6 months."

The following illustration shows the inputs and outputs of a risk mitigation strategy.

OCA risk mitigation strategy



The following table provides an example of a risk tracking tool.

Risk category	Severity	Probability	Risk description	Mitigation actions	Owner	Status	Due date
Resourcing	Medium	High	Security SME is taking a leave of absence that overlaps	Onboard and train backup security SME on specific tests and	Martha Rivera	In progress	31 March 2025

Risk category	Severity	Probability	Risk description	Mitigation actions	Owner	Status	Due date
			with our testing and cutover phase.	cutover planning.			

FAQ

Q. Why is the risk mitigation strategy valuable?

A. The risk mitigation strategy and plan offer a structured way to gain visibility into people-related issues that can stall, derail, or delay a cloud transformation. This process helps ensure that deliverables are on time, on budget, and produced with high quality, while offering an integrated approach to identify, assess, and address risks with the cloud transformation team.

Q. When should you use it?

A. Use a risk mitigation strategy and plan at the start of the program to design the format and establish risk dimensions. Review the strategy and plan on a regular cadence and update them as required.

Q. What types of issues are classified as a people-related and fall into the scope of this activity?

A. People-related issues are any non-technical issues that could hinder the cloud journey, such as leadership misalignment, prioritization differences, communication breakdowns, and cloud skill shortages.

Q. Who should be involved in this activity?

A. Participants should include the executive sponsor, cloud leader, OCA leader, HR lead, internal communications team, workstream leads, the project management office (PMO), and engagement managers.

Q. What are the inputs to this strategy and plan?

A. Inputs include culture assessment, [organization readiness assessment](#), leadership assessment review workshop, user readiness assessment, Migration Readiness Assessment (MRA) and Migration Readiness Planning (MRP), program risk log, and status reports.

Q. What are the outputs of this activity?

A. This activity produces the risk identification and management process and tracking tools that will be integrated into the overall program risk processes.

Q. Why should time be spent on this activity?

A. The risk mitigation strategy and plan ensure a seamless and integrated process to manage status, issues, and escalations, and to resolve conflicts before they block or slow down your cloud journey.

Additional steps

To develop a risk mitigation strategy and plan, follow these steps:

1. Collect potential people risks from various sources such as leadership alignment interviews, organizational readiness assessments, workstream leads, and status reports.
2. Evaluate and prioritize risks.
3. Assign risks to owners for mitigation and disposition.
4. Determine priorities for action, and assess the risks of not acting on identified issues.
5. Develop a risk escalation process.
6. Ensure that the people-related risk management process is integrated with customer and program processes.
7. Draft the risk mitigation strategy and plan.
8. Review and validate the strategy with the cloud leadership team.
9. Obtain customer and leadership signoff on the risk mitigation strategy and plan.
10. Conduct periodic risk management review meetings.
11. Track risk status on an ongoing basis.

By focusing on these elements and best practices, you can develop a comprehensive risk mitigation strategy for your organization that supports cloud transformation, addresses potential roadblocks, and ensures a smoother transition to the cloud environment.

3.6 Sponsor roadmap

Overview

Sponsor support and actions are critical levers for driving adoption. Having an active and visible sponsor is the most significant factor in achieving change adoption. The sponsor's active engagement and presence are instrumental in establishing the desired behaviors that are expected of individuals and organizations. A structured process secures consistency in messaging and helps achieve intended organizational objectives.

An effective sponsor roadmap delivers significant benefits:

- Accelerates cloud adoption through consistent leadership support
- Enhances alignment between leadership and transformation objectives
- Reduces resistance to change by demonstrating top-down commitment
- Improves communication effectiveness across the organization
- Increases the likelihood of achieving desired business outcomes
- Supports the cultural change that's necessary for successful cloud transformation

Best practices

To develop the sponsor roadmap, secure sponsor commitment at the beginning of the cloud program and take initial steps to:

- Provide general awareness and understanding of the sponsorship process to executives and cloud leaders.
- Onboard sponsors and provide role descriptions, accountability expectations, key messaging, and implementation timeline.
- Use key messaging from the [business case for change](#) to reinforce the cloud migration vision, benefits to the enterprise and stakeholder groups, and overall business value.
- Reinforce the message that the OCA team will be there every step of the way to foster commitment.

Consider both business and IT sponsors. When cloud adoption is a key component of your business strategy and outcomes, you must have sponsors from the business side of your organization, such as an executive sponsor and lines of business sponsors.

To design the sponsor roadmap:

- Review the strategic vision, business case, and outputs from earlier workshops and organizational analysis to gain insights into benefits and business value. Inputs often include:
 - Strategic vision
 - Business case
 - Discovery and other workshop findings and outputs
 - Executive sponsor and leadership interviews
 - Feedback from:
 - Change leadership stakeholder (If available)
 - Communications stakeholder (If available)
 - Training stakeholder (if available)
 - HR stakeholder
- Identify leaders and their assigned stakeholders who are essential for program messaging and engagement, will implement the program, and will engage with the selected stakeholders. At a minimum, this would include the executive sponsor, project leader, change leader, change agents or champions, internal change team liaison, internal communications, and HR.
- Define engagement objectives:
 - Understand the role of key stakeholders in implementing the program.
 - Define critical objectives for engaging selected stakeholders on a regular basis.
 - Socialize, discuss, and finalize engagement objectives with leaders who are responsible for achieving these objectives.
- Discuss the format and frequency of stakeholder engagements.
- Assess the engagement quality and address gaps. Establish a cadence for progress monitoring, review, and support.
 - Meet regularly with leaders to self-assess the current opinions of stakeholders against objectives.
 - Discuss the formats and frequency of stakeholder engagements. Is the current approach sufficient? Are new solutions required?

- Identify adjustments to the current approach and design new solutions for engaging stakeholders.
- Synthesize inputs and leader discussions to develop the OCA sponsor roadmap.

To implement the sponsor roadmap:

- Develop specific objectives for stakeholder and leadership action plans that align to the OCA sponsor roadmap.
- Engage with stakeholders as defined in the roadmap.
- Track the progress of planned actions to measure commitment risk.
- Update leadership action plans at appropriate intervals (at a minimum, quarterly) as phases and program risks change over time.

To be an effective sponsor:

- Stay focused on the vision, and remain active and visible to others throughout the project lifecycle.
- Communicate clearly and often, and provide targeted messaging to all stakeholder groups.
- Do not delegate sponsorship. Employees need to see ownership and accountability from their leaders.
- Lead from the front by clearly demonstrating your support for the project to empower your teams.
- Engage others in your business, and develop a sponsorship coalition to expand change ownership.
- Manage resistance by listening and responding to stakeholder feedback.
- Reinforce the change by rewarding and celebrating successes.
- Educate yourself on the people side of change and take a programmatic approach by applying the OCA 6-Point Framework. Be willing to devote the necessary amount of time and resources to address your sponsorship responsibilities.

Here are two examples of sponsor roadmaps, in the form of worksheets for leadership action plans.

Insert leader name | Sponsor/leadership action plan

Start month - end month | Team

		Month Year	Month Year	Month Year
Key initiative events and goals	• Insert key project events	• Insert key project events	• Insert key project events	• Insert key project events
PLAN	•	•	•	•
Leader names and action items	□ Target date Insert task	<i>Date complete</i>	□ Target date Insert task	<i>Date complete</i>
DO	□	□	□ Target date Insert task	<i>Date complete</i>
	□	□	□	□
Feedback: What worked well, what did not?	□ Feedback on activity/task	<i>Feedback owner</i>	□ Feedback on activity/task	<i>Feedback owner</i>
CHECK	□	□	□ Feedback on activity/task	<i>Feedback owner</i>
	□	□	□	□
Action items for team to maintain, update, or correct	□ Insert leader or elevate team follow-up activities	<i>Owner</i>	□ Insert leader or elevate team follow-up activities	<i>Owner</i>
ACT	□	□	□ Insert leader or elevate team follow-up activities	<i>Owner</i>
	□	□	□	□

Example action plan for achieving transformation goals

Name:		Position:		Date:	
Goal #1 What do I want to accomplish? <i>This should be aligned with one or more of your leader's goals.</i>					
Goal name	Metric target	Start date	End (due) date		
Reduce total cost of ownership (TCO) by 10 % by 202x a. Remain on target for transformation dates b. Maintain less than 2% delay	Annual TCO is reduced by 10%	202x	202x		

MY ACTIONS			
How can I accomplish goal #1?			
Goal name	Start	Finish	Status
1. Participate in planning.			25% = in progress
2. Allocate resources to transformation activities.			100%
3. Implement cost measurement plan.			50% = in progress
4. Communicate transformation goals.			100%
5. Ensure FinOps representation to evaluate TCO targets.			25% = in progress

FAQ

Q. When do you conduct this activity?

A. Secure sponsor commitment from the beginning of the cloud migration and modernization process. Provide general awareness and understanding of the sponsorship process to executives and cloud migration and modernization leaders. Onboard sponsors appropriately and provide role descriptions, accountability, key messaging, sponsor roadmap, and timeline. Reinforce the message that the change acceleration team will be involved in every step of the process to ensure commitment. Key messaging reinforces the cloud migration vision, benefits, and overall business value.

Q. Who should be involved in this activity?

A. Participants should include the executive sponsor, cloud leader, OCA leader, HR lead, the internal communications team, workstream leads, the project management office (PMO), and engagement managers.

Q. What are the inputs to this activity?

A. Inputs include the business case, discovery documentation findings, executive sponsor and HR interviews, stakeholder analysis, cloud strategy, and business value realization plans.

Q. What are the outputs of this activity?

A. Outputs include the identification of sponsorship needs, the sponsor plan design, and an implementation plan for sponsorship.

Additional steps

1. After you create your sponsor roadmap and implementation schedule, review it with workstream members and stakeholders, and refine it based on feedback.
2. As you roll it out, be prepared to modify it or iterate to align to program progress.
3. Consider asking a peer outside the cloud program to review the roadmap and ask them questions such as:
 - Is the sponsor roadmap easy to understand?
 - Can you explain it to others?
 - Does it address all leader stakeholder groups?
 - Is it achievable?
 - Where should it be shared? Which audiences need to hear about it?
 - Is it compelling enough to create a sense of urgency to change and adopt the cloud?
4. Track your progress, evaluate sponsor effectiveness by obtaining feedback from stakeholders, and adjust leadership action plans as needed.

By focusing on these elements and best practices, you can develop a comprehensive sponsor roadmap that supports cloud transformation, ensures consistent leadership engagement, and drives adoption throughout the organization.

3.7 Sustainability plan

Overview

The sustainability plan provides a proactive approach for transitioning organizational change acceleration activities from project status to business as usual (BAU). By developing this plan early, in the *Envision the Future* phase, you can create a roadmap for ensuring ownership of the cloud journey and minimize the risk of abandoning the project soon after its adoption.

Align planning activities with the cloud leadership team to understand and establish expectations beyond the initial phases of the cloud transformation journey. Consider:

- Organizational changes
- Gaps in positions, roles, and responsibilities
- Communication needs
- Additional training requirements
- Knowledge libraries or repositories
- Business metrics correlated with OCA measures

The sustainability plan evolves frequently; capture needs from status meetings, retrospectives, and risk, action, issue, dependency (RAID) logs throughout the cloud journey.

An effective sustainability plan delivers significant benefits:

- Ensures long-term adoption of cloud technologies and practices
- Reduces the risk of reverting to old ways of working
- Embeds cloud-centric culture into the organization's ways of working
- Maximizes return on investment (ROI) in cloud transformation efforts
- Supports continuous improvement and innovation
- Enhances organizational agility and adaptability

Best practices

In the early stages of the cloud journey, the move to the cloud is likely to be treated as a high-priority project or initiative. As such, the project often influences project and leadership behaviors and activities that contribute to success. For example:


Project behaviors and activities:

- Metrics
- Change acceleration and communications
- Assertive change agents and consultants

Leadership behaviors and activities:

- Creating important projects that the organization can support, and sustaining momentum
- Defining the long-term consequences of activity or non-activity for these projects
- Asking questions about the project and value realization
- Steering committee meetings

However, at some point, the journey to the cloud should stop being treated as a project or initiative and become business as usual.



Sustaining improvement gains over time requires a proactive and systematic approach for creating internal sustainability and ownership.

FAQ

Q. Why is the sustainability plan important?

A. The sustainability plan looks beyond the initial cloud migration phase to secure the necessary steps for adopting the future state model and standing the test of time. It provides a mechanism for future-proofing the cloud transformation as people and technologies evolve.

Q. What are the inputs to this plan?

A. Inputs include organizational change acceleration ownership vision and goals and any gaps that need to be addressed from an OCA perspective (for example, retrospectives, culture assessment, organization readiness assessment, leadership readiness review workshop, user readiness assessment, program risk log, and status reports).

Q. What are the outputs of this activity?

A. The main output is a project closure and sustainability plan that establishes ongoing ownership of OCA.

Q. How do we ensure that our CCoE is also sustainable?

A. As part of the AWS Cloud Center of Excellence (CCoE) methodology, we recommend that you evaluate the composition of your CCoE (the people, the processes, the tools, the policies, and so on) every quarter. As with the cloud transformation program, the CCoE will evolve as a leadership function.

Q. Where is sustainability applied in the cloud transformation lifecycle, and when should it be started?

A. Initiate the design of the sustainability plan in the *Envision the Future* stage of the cloud transformation lifecycle, and then scale the program.

Q. Why should time be spent on this activity?

A. Your cloud transformation journey will continue to introduce substantial changes to the organization as you move through migration, optimization, and innovation. Your ability to achieve and sustain desired business outcomes from your cloud program is directly proportional to your ability to own and sustain organizational change acceleration.

Additional steps

Here are the recommended steps to plan for ownership of future phases of your organizational acceleration strategy:

1. Review your cloud strategy and objectives. Are you on track to achieve desired business outcomes? Which people-related barriers do you need to overcome to achieve or accelerate your desired business outcomes?
2. Develop the vision and goals for organizational change acceleration ownership of the future state. Develop and validate assumptions.
3. If you have a Cloud Center of Excellence (CCoE), evaluate your Cloud Business Office (CBO) capability maturity as it relates to leadership, change acceleration, training, and communications.
4. Review your change planning materials and feedback. For example:
 - Which key risks have been identified?
 - What unexpected feedback has emerged?
 - Where do you see natural next steps for the organization?
5. Conduct frequent retrospectives on the change acceleration workstream. Solicit input from all other workstreams. What is working well? What can be improved?
6. Evaluate your organization's ability to own and sustain each major component of your organizational change acceleration strategy:
 - Business value and outcome realization tracking
 - Whether leadership is driving cloud adoption
 - Whether executive sponsors are communicating the case for change and removing blockers
 - Culture change for achieving optimal benefits of the cloud
 - Communications to affected stakeholders
 - Training plans for current and future cloud needs
 - Talent acquisition strategies that are aligned to future cloud needs
 - Talent management and workforce transformation that is aligned to the cloud strategy
 - Strategy to retain valuable cloud talent after they become digitally fluent and skilled on cloud solutions, or have obtained cloud certifications
7. For each component, consider resources, competencies, processes, structure, behaviors and tasks, and consequences and incentive systems.

Focus area	Considerations
Resources	<ul style="list-style-type: none">• How many resources were required to lead and implement the current (cloud or on-premises) scope?• How many resources will be needed to lead and implement the future state?• How can we close the gap and create more ownership?
Competencies	<ul style="list-style-type: none">• What is the current competency compared with the desired level of competency?• What is the highest priority (for example, planning, training, communications)?• How can we close the gap and increase competency through formal or informal training, shadowing, and experiential activities?• How do we measure or ensure proficiency?
Processes	<ul style="list-style-type: none">• What is the current process for implementing the organizational change acceleration component? Where are the likely failure points?• Is there a future state process that has less friction and can be made more adaptable by frequent feedback loops? How can it be more automated?• Who is accountable for the process? Does a responsible, accountable, consulted, informed (RACI) matrix exist?

Focus area	Considerations
Structure	<ul style="list-style-type: none"> • Do we have the right structure (for example, centralized, distributed, or embedded) to support our future state? • Do we have people in the right locations to support the organizational change acceleration strategy? • Do we have the appropriate lines of accountability and feedback?
Behaviors and tasks	<ul style="list-style-type: none"> • What are the behaviors or tasks needed to implement the organizational change acceleration component? • What are the guiding principles of organizational change acceleration? • Are the behaviors or tasks clearly defined so that others can implement them?
Consequences and incentive systems	<ul style="list-style-type: none"> • Are formal and informal incentive systems aligned to support the desired future state behaviors? • What additional positive consequences or incentives can we put into place to encourage future state behaviors? • Are desired behaviors being met with negative consequences, inadvertently? • Are undesired behaviors being rewarded, inadvertently?

By focusing on these elements and best practices, you can develop a comprehensive sustainability plan that ensures long-term success in your organization's cloud transformation journey.

Resources

References

- [Accelerating your return on cloud investment by adopting a strategic transformation and change methodology](#)
- [AWS Change Acceleration 6-Point Framework and Organizational Change Management Toolkit](#)
- [AWS Organizational Change Acceleration \(OCA\) 6-Point Framework – 1. Mobilize Team](#)
- [AWS Organizational Change Acceleration \(OCA\) 6-Point Framework – 2. Align Leaders](#)
- [AWS Organizational Change Acceleration \(OCA\) 6-Point Framework – 4. Engage the Organization](#)
- [AWS Organizational Change Acceleration \(OCA\) 6-Point Framework – 5. Enable Capacity](#)
- [AWS Organizational Change Acceleration \(OCA\) 6-Point Framework – 6. Make Culture Change Stick](#)
- [AWS Cloud Adoption Framework \(CAF\)](#)
- [AWS Cloud Adoption Framework \(CAF\) People Perspective](#)
- [2024 IT Skills and Salary \(Skillsoft Global Knowledge report\)](#)
- [AWS Certification and Training](#)
- [AWS GameDay](#)
- [AWS Solutions-Focused Immersion Days](#)
- [AWS Learning Needs Analysis \(LNA\)](#)

Partners

- **Accenture**
 - [Contact Partner](#)
 - [Contact the Accenture AWS Business Group](#)
 - [Future Talent Platform](#)
 - [Accenture and AWS take you further faster](#)
- **Deloitte**
 - [Contact Partner](#)

- [AWS and Deloitte](#)
- [Where Innovation Meets Impact](#)
- **PwC**
 - [Contact Partner](#)
 - [PwC and AWS](#)
- **Slalom**
 - [Contact Partner](#)
 - [AWS and Slalom Launch Centers](#)
- **Roberts Group Consulting**
 - [Contact Partner](#)

Contributors

- Melanie Gladwell, AWS Senior Practice Manager
- Scott Watson, AWS People Transformation Lead
- Tierra Jennings-Hill, AWS People Transformation Lead
- Nicole Lenz, AWS Sales Transformation Lead
- Jermel Moody, AWS Change Acceleration Lead
- Travis McNeal, AWS Change Acceleration Lead

Document history

The following table describes significant changes to this guide. If you want to be notified about future updates, you can subscribe to an [RSS feed](#).

Change	Description	Date
Initial publication	—	January 31, 2025

AWS Prescriptive Guidance glossary

The following are commonly used terms in strategies, guides, and patterns provided by AWS Prescriptive Guidance. To suggest entries, please use the **Provide feedback** link at the end of the glossary.

Numbers

7 Rs

Seven common migration strategies for moving applications to the cloud. These strategies build upon the 5 Rs that Gartner identified in 2011 and consist of the following:

- **Refactor/re-architect** – Move an application and modify its architecture by taking full advantage of cloud-native features to improve agility, performance, and scalability. This typically involves porting the operating system and database. Example: Migrate your on-premises Oracle database to the Amazon Aurora PostgreSQL-Compatible Edition.
- **Replatform (lift and reshape)** – Move an application to the cloud, and introduce some level of optimization to take advantage of cloud capabilities. Example: Migrate your on-premises Oracle database to Amazon Relational Database Service (Amazon RDS) for Oracle in the AWS Cloud.
- **Repurchase (drop and shop)** – Switch to a different product, typically by moving from a traditional license to a SaaS model. Example: Migrate your customer relationship management (CRM) system to Salesforce.com.
- **Rehost (lift and shift)** – Move an application to the cloud without making any changes to take advantage of cloud capabilities. Example: Migrate your on-premises Oracle database to Oracle on an EC2 instance in the AWS Cloud.
- **Relocate (hypervisor-level lift and shift)** – Move infrastructure to the cloud without purchasing new hardware, rewriting applications, or modifying your existing operations. You migrate servers from an on-premises platform to a cloud service for the same platform. Example: Migrate a Microsoft Hyper-V application to AWS.
- **Retain (revisit)** – Keep applications in your source environment. These might include applications that require major refactoring, and you want to postpone that work until a later time, and legacy applications that you want to retain, because there's no business justification for migrating them.

- **Retire** – Decommission or remove applications that are no longer needed in your source environment.

A

A2A (Agent-to-Agent)

A stateful protocol for agent-to-agent collaboration supporting task delegation and state transfer.

ABAC

See [attribute-based access control](#).

abstracted services

See [managed services](#).

ACID

See [atomicity, consistency, isolation, durability](#).

active-active migration

A database migration method in which the source and target databases are kept in sync (by using a bidirectional replication tool or dual write operations), and both databases handle transactions from connecting applications during migration. This method supports migration in small, controlled batches instead of requiring a one-time cutover. It's more flexible but requires more work than [active-passive migration](#).

active-passive migration

A database migration method in which the source and target databases are kept in sync, but only the source database handles transactions from connecting applications while data is replicated to the target database. The target database doesn't accept any transactions during migration.

Agent

An AI system that can autonomously reason, plan, and take actions using tools to achieve goals.

Agent Ops

Operational practices for building, testing, deploying, and running AI agents in production at scale.

aggregate function

A SQL function that operates on a group of rows and calculates a single return value for the group. Examples of aggregate functions include SUM and MAX.

AI

See [artificial intelligence](#).

AIOps

See [artificial intelligence operations](#).

anonymization

The process of permanently deleting personal information in a dataset. Anonymization can help protect personal privacy. Anonymized data is no longer considered to be personal data.

anti-pattern

A frequently used solution for a recurring issue where the solution is counter-productive, ineffective, or less effective than an alternative.

application control

A security approach that allows the use of only approved applications in order to help protect a system from malware.

application portfolio

A collection of detailed information about each application used by an organization, including the cost to build and maintain the application, and its business value. This information is key to [the portfolio discovery and analysis process](#) and helps identify and prioritize the applications to be migrated, modernized, and optimized.

artificial intelligence (AI)

The field of computer science that is dedicated to using computing technologies to perform cognitive functions that are typically associated with humans, such as learning, solving problems, and recognizing patterns. For more information, see [What is Artificial Intelligence?](#)

artificial intelligence operations (AIOps)

The process of using machine learning techniques to solve operational problems, reduce operational incidents and human intervention, and increase service quality. For more information about how AIOps is used in the AWS migration strategy, see the [operations integration guide](#).

asymmetric encryption

An encryption algorithm that uses a pair of keys, a public key for encryption and a private key for decryption. You can share the public key because it isn't used for decryption, but access to the private key should be highly restricted.

atomicity, consistency, isolation, durability (ACID)

A set of software properties that guarantee the data validity and operational reliability of a database, even in the case of errors, power failures, or other problems.

attribute-based access control (ABAC)

The practice of creating fine-grained permissions based on user attributes, such as department, job role, and team name. For more information, see [ABAC for AWS](#) in the AWS Identity and Access Management (IAM) documentation.

authoritative data source

A location where you store the primary version of data, which is considered to be the most reliable source of information. You can copy data from the authoritative data source to other locations for the purposes of processing or modifying the data, such as anonymizing, redacting, or pseudonymizing it.

Availability Zone

A distinct location within an AWS Region that is insulated from failures in other Availability Zones and provides inexpensive, low-latency network connectivity to other Availability Zones in the same Region.

AWS Cloud Adoption Framework (AWS CAF)

A framework of guidelines and best practices from AWS to help organizations develop an efficient and effective plan to move successfully to the cloud. AWS CAF organizes guidance into six focus areas called perspectives: business, people, governance, platform, security, and operations. The business, people, and governance perspectives focus on business skills and processes; the platform, security, and operations perspectives focus on technical skills and processes. For example, the people perspective targets stakeholders who handle human resources (HR), staffing functions, and people management. For this perspective, AWS CAF provides guidance for people development, training, and communications to help ready the organization for successful cloud adoption. For more information, see the [AWS CAF website](#) and the [AWS CAF whitepaper](#).

AWS Workload Qualification Framework (AWS WQF)

A tool that evaluates database migration workloads, recommends migration strategies, and provides work estimates. AWS WQF is included with AWS Schema Conversion Tool (AWS SCT). It analyzes database schemas and code objects, application code, dependencies, and performance characteristics, and provides assessment reports.

B

bad bot

A [bot](#) that is intended to disrupt or cause harm to individuals or organizations.

BCP

See [business continuity planning](#).

behavior graph

A unified, interactive view of resource behavior and interactions over time. You can use a behavior graph with Amazon Detective to examine failed logon attempts, suspicious API calls, and similar actions. For more information, see [Data in a behavior graph](#) in the Detective documentation.

big-endian system

A system that stores the most significant byte first. See also [endianness](#).

binary classification

A process that predicts a binary outcome (one of two possible classes). For example, your ML model might need to predict problems such as "Is this email spam or not spam?" or "Is this product a book or a car?"

bloom filter

A probabilistic, memory-efficient data structure that is used to test whether an element is a member of a set.

blue/green deployment

A deployment strategy where you create two separate but identical environments. You run the current application version in one environment (blue) and the new application version in the other environment (green). This strategy helps you quickly roll back with minimal impact.

bot

A software application that runs automated tasks over the internet and simulates human activity or interaction. Some bots are useful or beneficial, such as web crawlers that index information on the internet. Some other bots, known as *bad bots*, are intended to disrupt or cause harm to individuals or organizations.

botnet

Networks of [bots](#) that are infected by [malware](#) and are under the control of a single party, known as a *bot herder* or *bot operator*. Botnets are the best-known mechanism to scale bots and their impact.

branch

A contained area of a code repository. The first branch created in a repository is the *main branch*. You can create a new branch from an existing branch, and you can then develop features or fix bugs in the new branch. A branch you create to build a feature is commonly referred to as a *feature branch*. When the feature is ready for release, you merge the feature branch back into the main branch. For more information, see [About branches](#) (GitHub documentation).

break-glass access

In exceptional circumstances and through an approved process, a quick means for a user to gain access to an AWS account that they don't typically have permissions to access. For more information, see the [Implement break-glass procedures](#) indicator in the AWS Well-Architected guidance.

brownfield strategy

The existing infrastructure in your environment. When adopting a brownfield strategy for a system architecture, you design the architecture around the constraints of the current systems and infrastructure. If you are expanding the existing infrastructure, you might blend brownfield and [greenfield](#) strategies.

buffer cache

The memory area where the most frequently accessed data is stored.

business capability

What a business does to generate value (for example, sales, customer service, or marketing). Microservices architectures and development decisions can be driven by business capabilities.

For more information, see the [Organized around business capabilities](#) section of the [Running containerized microservices on AWS](#) whitepaper.

business continuity planning (BCP)

A plan that addresses the potential impact of a disruptive event, such as a large-scale migration, on operations and enables a business to resume operations quickly.

C

CAF

See [AWS Cloud Adoption Framework](#).

canary deployment

The slow and incremental release of a version to end users. When you are confident, you deploy the new version and replace the current version in its entirety.

CCoE

See [Cloud Center of Excellence](#).

CDC

See [change data capture](#).

change data capture (CDC)

The process of tracking changes to a data source, such as a database table, and recording metadata about the change. You can use CDC for various purposes, such as auditing or replicating changes in a target system to maintain synchronization.

chaos engineering

Intentionally introducing failures or disruptive events to test a system's resilience. You can use [AWS Fault Injection Service \(AWS FIS\)](#) to perform experiments that stress your AWS workloads and evaluate their response.

CI/CD

See [continuous integration and continuous delivery](#).

classification

A categorization process that helps generate predictions. ML models for classification problems predict a discrete value. Discrete values are always distinct from one another. For example, a model might need to evaluate whether or not there is a car in an image.

Citizen Developer

A business user who creates AI applications using no-code/low-code platforms without specialized technical skills.

client-side encryption

Encryption of data locally, before the target AWS service receives it.

Cloud Center of Excellence (CCoE)

A multi-disciplinary team that drives cloud adoption efforts across an organization, including developing cloud best practices, mobilizing resources, establishing migration timelines, and leading the organization through large-scale transformations. For more information, see the [CCoE posts](#) on the AWS Cloud Enterprise Strategy Blog.

cloud computing

The cloud technology that is typically used for remote data storage and IoT device management. Cloud computing is commonly connected to [edge computing](#) technology.

cloud operating model

In an IT organization, the operating model that is used to build, mature, and optimize one or more cloud environments. For more information, see [Building your Cloud Operating Model](#).

cloud stages of adoption

The four phases that organizations typically go through when they migrate to the AWS Cloud:

- Project – Running a few cloud-related projects for proof of concept and learning purposes
- Foundation – Making foundational investments to scale your cloud adoption (e.g., creating a landing zone, defining a CCoE, establishing an operations model)
- Migration – Migrating individual applications
- Re-invention – Optimizing products and services, and innovating in the cloud

These stages were defined by Stephen Orban in the blog post [The Journey Toward Cloud-First & the Stages of Adoption](#) on the AWS Cloud Enterprise Strategy blog. For information about how they relate to the AWS migration strategy, see the [migration readiness guide](#).

CMDB

See [configuration management database](#).

code repository

A location where source code and other assets, such as documentation, samples, and scripts, are stored and updated through version control processes. Common cloud repositories include GitHub or Bitbucket Cloud. Each version of the code is called a *branch*. In a microservice structure, each repository is devoted to a single piece of functionality. A single CI/CD pipeline can use multiple repositories.

cold cache

A buffer cache that is empty, not well populated, or contains stale or irrelevant data. This affects performance because the database instance must read from the main memory or disk, which is slower than reading from the buffer cache.

cold data

Data that is rarely accessed and is typically historical. When querying this kind of data, slow queries are typically acceptable. Moving this data to lower-performing and less expensive storage tiers or classes can reduce costs.

computer vision (CV)

A field of [AI](#) that uses machine learning to analyze and extract information from visual formats such as digital images and videos. For example, Amazon SageMaker AI provides image processing algorithms for CV.

configuration drift

For a workload, a configuration change from the expected state. It might cause the workload to become noncompliant, and it's typically gradual and unintentional.

configuration management database (CMDB)

A repository that stores and manages information about a database and its IT environment, including both hardware and software components and their configurations. You typically use data from a CMDB in the portfolio discovery and analysis stage of migration.

conformance pack

A collection of AWS Config rules and remediation actions that you can assemble to customize your compliance and security checks. You can deploy a conformance pack as a single entity in

an AWS account and Region, or across an organization, by using a YAML template. For more information, see [Conformance packs](#) in the AWS Config documentation.

continuous integration and continuous delivery (CI/CD)

The process of automating the source, build, test, staging, and production stages of the software release process. CI/CD is commonly described as a pipeline. CI/CD can help you automate processes, improve productivity, improve code quality, and deliver faster. For more information, see [Benefits of continuous delivery](#). CD can also stand for *continuous deployment*. For more information, see [Continuous Delivery vs. Continuous Deployment](#).

CV

See [computer vision](#).

D

data at rest

Data that is stationary in your network, such as data that is in storage.

data classification

A process for identifying and categorizing the data in your network based on its criticality and sensitivity. It is a critical component of any cybersecurity risk management strategy because it helps you determine the appropriate protection and retention controls for the data. Data classification is a component of the security pillar in the AWS Well-Architected Framework. For more information, see [Data classification](#).

data drift

A meaningful variation between the production data and the data that was used to train an ML model, or a meaningful change in the input data over time. Data drift can reduce the overall quality, accuracy, and fairness in ML model predictions.

data in transit

Data that is actively moving through your network, such as between network resources.

data mesh

An architectural framework that provides distributed, decentralized data ownership with centralized management and governance.

data minimization

The principle of collecting and processing only the data that is strictly necessary. Practicing data minimization in the AWS Cloud can reduce privacy risks, costs, and your analytics carbon footprint.

data perimeter

A set of preventive guardrails in your AWS environment that help make sure that only trusted identities are accessing trusted resources from expected networks. For more information, see [Building a data perimeter on AWS](#).

data preprocessing

To transform raw data into a format that is easily parsed by your ML model. Preprocessing data can mean removing certain columns or rows and addressing missing, inconsistent, or duplicate values.

data provenance

The process of tracking the origin and history of data throughout its lifecycle, such as how the data was generated, transmitted, and stored.

data subject

An individual whose data is being collected and processed.

data warehouse

A data management system that supports business intelligence, such as analytics. Data warehouses commonly contain large amounts of historical data, and they are typically used for queries and analysis.

database definition language (DDL)

Statements or commands for creating or modifying the structure of tables and objects in a database.

database manipulation language (DML)

Statements or commands for modifying (inserting, updating, and deleting) information in a database.

DDL

See [database definition language](#).

deep ensemble

To combine multiple deep learning models for prediction. You can use deep ensembles to obtain a more accurate prediction or for estimating uncertainty in predictions.

deep learning

An ML subfield that uses multiple layers of artificial neural networks to identify mapping between input data and target variables of interest.

defense-in-depth

An information security approach in which a series of security mechanisms and controls are thoughtfully layered throughout a computer network to protect the confidentiality, integrity, and availability of the network and the data within. When you adopt this strategy on AWS, you add multiple controls at different layers of the AWS Organizations structure to help secure resources. For example, a defense-in-depth approach might combine multi-factor authentication, network segmentation, and encryption.

delegated administrator

In AWS Organizations, a compatible service can register an AWS member account to administer the organization's accounts and manage permissions for that service. This account is called the *delegated administrator* for that service. For more information and a list of compatible services, see [Services that work with AWS Organizations](#) in the AWS Organizations documentation.

deployment

The process of making an application, new features, or code fixes available in the target environment. Deployment involves implementing changes in a code base and then building and running that code base in the application's environments.

development environment

See [environment](#).

detective control

A security control that is designed to detect, log, and alert after an event has occurred. These controls are a second line of defense, alerting you to security events that bypassed the preventative controls in place. For more information, see [Detective controls](#) in *Implementing security controls on AWS*.

development value stream mapping (DVSM)

A process used to identify and prioritize constraints that adversely affect speed and quality in a software development lifecycle. DVSM extends the value stream mapping process originally designed for lean manufacturing practices. It focuses on the steps and teams required to create and move value through the software development process.

digital twin

A virtual representation of a real-world system, such as a building, factory, industrial equipment, or production line. Digital twins support predictive maintenance, remote monitoring, and production optimization.

dimension table

In a [star schema](#), a smaller table that contains data attributes about quantitative data in a fact table. Dimension table attributes are typically text fields or discrete numbers that behave like text. These attributes are commonly used for query constraining, filtering, and result set labeling.

disaster

An event that prevents a workload or system from fulfilling its business objectives in its primary deployed location. These events can be natural disasters, technical failures, or the result of human actions, such as unintentional misconfiguration or a malware attack.

disaster recovery (DR)

The strategy and process you use to minimize downtime and data loss caused by a [disaster](#). For more information, see [Disaster Recovery of Workloads on AWS: Recovery in the Cloud](#) in the AWS Well-Architected Framework.

DML

See [database manipulation language](#).

domain-driven design

An approach to developing a complex software system by connecting its components to evolving domains, or core business goals, that each component serves. This concept was introduced by Eric Evans in his book, *Domain-Driven Design: Tackling Complexity in the Heart of Software* (Boston: Addison-Wesley Professional, 2003). For information about how you can use domain-driven design with the strangler fig pattern, see [Modernizing legacy Microsoft ASP.NET \(ASMX\) web services incrementally by using containers and Amazon API Gateway](#).

DR

See [disaster recovery](#).

drift detection

Tracking deviations from a baselined configuration. For example, you can use AWS CloudFormation to [detect drift in system resources](#), or you can use AWS Control Tower to [detect changes in your landing zone](#) that might affect compliance with governance requirements.

DVSM

See [development value stream mapping](#).

E

EDA

See [exploratory data analysis](#).

EDI

See [electronic data interchange](#).

edge computing

The technology that increases the computing power for smart devices at the edges of an IoT network. When compared with [cloud computing](#), edge computing can reduce communication latency and improve response time.

electronic data interchange (EDI)

The automated exchange of business documents between organizations. For more information, see [What is Electronic Data Interchange](#).

encryption

A computing process that transforms plaintext data, which is human-readable, into ciphertext.

encryption key

A cryptographic string of randomized bits that is generated by an encryption algorithm. Keys can vary in length, and each key is designed to be unpredictable and unique.

endianness

The order in which bytes are stored in computer memory. Big-endian systems store the most significant byte first. Little-endian systems store the least significant byte first.

endpoint

See [service endpoint](#).

endpoint service

A service that you can host in a virtual private cloud (VPC) to share with other users. You can create an endpoint service with AWS PrivateLink and grant permissions to other AWS accounts or to AWS Identity and Access Management (IAM) principals. These accounts or principals can connect to your endpoint service privately by creating interface VPC endpoints. For more information, see [Create an endpoint service](#) in the Amazon Virtual Private Cloud (Amazon VPC) documentation.

enterprise resource planning (ERP)

A system that automates and manages key business processes (such as accounting, [MES](#), and project management) for an enterprise.

envelope encryption

The process of encrypting an encryption key with another encryption key. For more information, see [Envelope encryption](#) in the AWS Key Management Service (AWS KMS) documentation.

environment

An instance of a running application. The following are common types of environments in cloud computing:

- development environment – An instance of a running application that is available only to the core team responsible for maintaining the application. Development environments are used to test changes before promoting them to upper environments. This type of environment is sometimes referred to as a *test environment*.
- lower environments – All development environments for an application, such as those used for initial builds and tests.
- production environment – An instance of a running application that end users can access. In a CI/CD pipeline, the production environment is the last deployment environment.

- upper environments – All environments that can be accessed by users other than the core development team. This can include a production environment, preproduction environments, and environments for user acceptance testing.

epic

In agile methodologies, functional categories that help organize and prioritize your work. Epics provide a high-level description of requirements and implementation tasks. For example, AWS CAF security epics include identity and access management, detective controls, infrastructure security, data protection, and incident response. For more information about epics in the AWS migration strategy, see the [program implementation guide](#).

ERP

See [enterprise resource planning](#).

exploratory data analysis (EDA)

The process of analyzing a dataset to understand its main characteristics. You collect or aggregate data and then perform initial investigations to find patterns, detect anomalies, and check assumptions. EDA is performed by calculating summary statistics and creating data visualizations.

F

fact table

The central table in a [star schema](#). It stores quantitative data about business operations. Typically, a fact table contains two types of columns: those that contain measures and those that contain a foreign key to a dimension table.

fail fast

A philosophy that uses frequent and incremental testing to reduce the development lifecycle. It is a critical part of an agile approach.

fault isolation boundary

In the AWS Cloud, a boundary such as an Availability Zone, AWS Region, control plane, or data plane that limits the effect of a failure and helps improve the resilience of workloads. For more information, see [AWS Fault Isolation Boundaries](#).

feature branch

See [branch](#).

features

The input data that you use to make a prediction. For example, in a manufacturing context, features could be images that are periodically captured from the manufacturing line.

feature importance

How significant a feature is for a model's predictions. This is usually expressed as a numerical score that can be calculated through various techniques, such as Shapley Additive Explanations (SHAP) and integrated gradients. For more information, see [Machine learning model interpretability with AWS](#).

feature transformation

To optimize data for the ML process, including enriching data with additional sources, scaling values, or extracting multiple sets of information from a single data field. This enables the ML model to benefit from the data. For example, if you break down the "2021-05-27 00:15:37" date into "2021", "May", "Thu", and "15", you can help the learning algorithm learn nuanced patterns associated with different data components.

few-shot prompting

Providing an [LLM](#) with a small number of examples that demonstrate the task and desired output before asking it to perform a similar task. This technique is an application of in-context learning, where models learn from examples (*shots*) that are embedded in prompts. Few-shot prompting can be effective for tasks that require specific formatting, reasoning, or domain knowledge. See also [zero-shot prompting](#).

FGAC

See [fine-grained access control](#).

fine-grained access control (FGAC)

The use of multiple conditions to allow or deny an access request.

flash-cut migration

A database migration method that uses continuous data replication through [change data capture](#) to migrate data in the shortest time possible, instead of using a phased approach. The objective is to keep downtime to a minimum.

FM

See [foundation model](#).

foundation model (FM)

A large deep-learning neural network that has been training on massive datasets of generalized and unlabeled data. FMs are capable of performing a wide variety of general tasks, such as understanding language, generating text and images, and conversing in natural language. For more information, see [What are Foundation Models](#).

FM gateway

A centralized intermediary that controls and normalizes access to [foundation models](#). Also known as an *LLM gateway*.

G

generative AI

A subset of [AI](#) models that have been trained on large amounts of data and that can use a simple text prompt to create new content and artifacts, such as images, videos, text, and audio. For more information, see [What is Generative AI](#).

geo blocking

See [geographic restrictions](#).

geographic restrictions (geo blocking)

In Amazon CloudFront, an option to prevent users in specific countries from accessing content distributions. You can use an allow list or block list to specify approved and banned countries. For more information, see [Restricting the geographic distribution of your content](#) in the CloudFront documentation.

Gitflow workflow

An approach in which lower and upper environments use different branches in a source code repository. The Gitflow workflow is considered legacy, and the [trunk-based workflow](#) is the modern, preferred approach.

golden image

A snapshot of a system or software that is used as a template to deploy new instances of that system or software. For example, in manufacturing, a golden image can be used to provision software on multiple devices and helps improve speed, scalability, and productivity in device manufacturing operations.

greenfield strategy

The absence of existing infrastructure in a new environment. When adopting a greenfield strategy for a system architecture, you can select all new technologies without the restriction of compatibility with existing infrastructure, also known as [brownfield](#). If you are expanding the existing infrastructure, you might blend brownfield and greenfield strategies.

guardrail

A high-level rule that helps govern resources, policies, and compliance across organizational units (OUs). *Preventive guardrails* enforce policies to ensure alignment to compliance standards. They are implemented by using service control policies and IAM permissions boundaries. *Detective guardrails* detect policy violations and compliance issues, and generate alerts for remediation. They are implemented by using AWS Config, AWS Security Hub CSPM, Amazon GuardDuty, AWS Trusted Advisor, Amazon Inspector, and custom AWS Lambda checks.

guardrails (AI)

Safety mechanisms that filter, validate, and constrain [agent](#) inputs and outputs to help ensure responsible and safe AI behavior.

H

HA

See [high availability](#).

heterogeneous database migration

Migrating your source database to a target database that uses a different database engine (for example, Oracle to Amazon Aurora). Heterogeneous migration is typically part of a re-architecting effort, and converting the schema can be a complex task. [AWS provides AWS SCT](#) that helps with schema conversions.

high availability (HA)

The ability of a workload to operate continuously, without intervention, in the event of challenges or disasters. HA systems are designed to automatically fail over, consistently deliver high-quality performance, and handle different loads and failures with minimal performance impact.

historian modernization

An approach used to modernize and upgrade operational technology (OT) systems to better serve the needs of the manufacturing industry. A *historian* is a type of database that is used to collect and store data from various sources in a factory.

holdout data

A portion of historical, labeled data that is withheld from a dataset that is used to train a [machine learning](#) model. You can use holdout data to evaluate the model performance by comparing the model predictions against the holdout data.

human-in-the-loop (HitL)

A workflow pattern where [agent](#) execution pauses for human review and approval at critical decision points.

homogeneous database migration

Migrating your source database to a target database that shares the same database engine (for example, Microsoft SQL Server to Amazon RDS for SQL Server). Homogeneous migration is typically part of a rehosting or replatforming effort. You can use native database utilities to migrate the schema.

hot data

Data that is frequently accessed, such as real-time data or recent translational data. This data typically requires a high-performance storage tier or class to provide fast query responses.

hotfix

An urgent fix for a critical issue in a production environment. Due to its urgency, a hotfix is usually made outside of the typical DevOps release workflow.

hypercare period

Immediately following cutover, the period of time when a migration team manages and monitors the migrated applications in the cloud in order to address any issues. Typically, this

period is 1–4 days in length. At the end of the hypercare period, the migration team typically transfers responsibility for the applications to the cloud operations team.

I

IaC

See [infrastructure as code](#).

identity-based policy

A policy attached to one or more IAM principals that defines their permissions within the AWS Cloud environment.

idle application

An application that has an average CPU and memory usage between 5 and 20 percent over a period of 90 days. In a migration project, it is common to retire these applications or retain them on premises.

IIoT

See [industrial Internet of Things](#).

immutable infrastructure

A model that deploys new infrastructure for production workloads instead of updating, patching, or modifying the existing infrastructure. Immutable infrastructures are inherently more consistent, reliable, and predictable than [mutable infrastructure](#). For more information, see the [Deploy using immutable infrastructure](#) best practice in the AWS Well-Architected Framework.

inbound (ingress) VPC

In an AWS multi-account architecture, a VPC that accepts, inspects, and routes network connections from outside an application. The [AWS Security Reference Architecture](#) recommends setting up your Network account with inbound, outbound, and inspection VPCs to protect the two-way interface between your application and the broader internet.

incremental migration

A cutover strategy in which you migrate your application in small parts instead of performing a single, full cutover. For example, you might move only a few microservices or users to the new system initially. After you verify that everything is working properly, you can incrementally

move additional microservices or users until you can decommission your legacy system. This strategy reduces the risks associated with large migrations.

Industry 4.0

A term that was introduced by [Klaus Schwab](#) in 2016 to refer to the modernization of manufacturing processes through advances in connectivity, real-time data, automation, analytics, and AI/ML.

infrastructure

All of the resources and assets contained within an application's environment.

infrastructure as code (IaC)

The process of provisioning and managing an application's infrastructure through a set of configuration files. IaC is designed to help you centralize infrastructure management, standardize resources, and scale quickly so that new environments are repeatable, reliable, and consistent.

industrial Internet of Things (IIoT)

The use of internet-connected sensors and devices in the industrial sectors, such as manufacturing, energy, automotive, healthcare, life sciences, and agriculture. For more information, see [Building an industrial Internet of Things \(IIoT\) digital transformation strategy](#).

inspection VPC

In an AWS multi-account architecture, a centralized VPC that manages inspections of network traffic between VPCs (in the same or different AWS Regions), the internet, and on-premises networks. The [AWS Security Reference Architecture](#) recommends setting up your Network account with inbound, outbound, and inspection VPCs to protect the two-way interface between your application and the broader internet.

Internet of Things (IoT)

The network of connected physical objects with embedded sensors or processors that communicate with other devices and systems through the internet or over a local communication network. For more information, see [What is IoT?](#)

interpretability

A characteristic of a machine learning model that describes the degree to which a human can understand how the model's predictions depend on its inputs. For more information, see [Machine learning model interpretability with AWS](#).

IoT

See [Internet of Things](#).

IT information library (ITIL)

A set of best practices for delivering IT services and aligning these services with business requirements. ITIL provides the foundation for ITSM.

IT service management (ITSM)

Activities associated with designing, implementing, managing, and supporting IT services for an organization. For information about integrating cloud operations with ITSM tools, see the [operations integration guide](#).

ITIL

See [IT information library](#).

ITSM

See [IT service management](#).

L

label-based access control (LBAC)

An implementation of mandatory access control (MAC) where the users and the data itself are each explicitly assigned a security label value. The intersection between the user security label and data security label determines which rows and columns can be seen by the user.

landing zone

A landing zone is a well-architected, multi-account AWS environment that is scalable and secure. This is a starting point from which your organizations can quickly launch and deploy workloads and applications with confidence in their security and infrastructure environment. For more information about landing zones, see [Setting up a secure and scalable multi-account AWS environment](#).

large language model (LLM)

A deep learning [AI](#) model that is pretrained on a vast amount of data. An LLM can perform multiple tasks, such as answering questions, summarizing documents, translating text into other languages, and completing sentences. For more information, see [What are LLMs](#).

large migration

A migration of 300 or more servers.

LBAC

See [label-based access control](#).

least privilege

The security best practice of granting the minimum permissions required to perform a task. For more information, see [Apply least-privilege permissions](#) in the IAM documentation.

lift and shift

See [7 Rs](#).

little-endian system

A system that stores the least significant byte first. See also [endianness](#).

LLM

See [large language model](#).

lower environments

See [environment](#).

M

machine learning (ML)

A type of artificial intelligence that uses algorithms and techniques for pattern recognition and learning. ML analyzes and learns from recorded data, such as Internet of Things (IoT) data, to generate a statistical model based on patterns. For more information, see [Machine Learning](#).

main branch

See [branch](#).

malware

Software that is designed to compromise computer security or privacy. Malware might disrupt computer systems, leak sensitive information, or gain unauthorized access. Examples of malware include viruses, worms, ransomware, Trojan horses, spyware, and keyloggers.

managed services

AWS services for which AWS operates the infrastructure layer, the operating system, and platforms, and you access the endpoints to store and retrieve data. Amazon Simple Storage Service (Amazon S3) and Amazon DynamoDB are examples of managed services. These are also known as *abstracted services*.

manufacturing execution system (MES)

A software system for tracking, monitoring, documenting, and controlling production processes that convert raw materials to finished products on the shop floor.

MAP

See [Migration Acceleration Program](#).

MCP

See [Model Context Protocol](#).

Model Context Protocol (MCP)

A stateless protocol for [agent](#)-to-[tool](#) communication.

MCP server

A service that exposes one or more [tools](#) through the [Model Context Protocol](#).

mechanism

A complete process in which you create a tool, drive adoption of the tool, and then inspect the results in order to make adjustments. A mechanism is a cycle that reinforces and improves itself as it operates. For more information, see [Building mechanisms](#) in the AWS Well-Architected Framework.

member account

All AWS accounts other than the management account that are part of an organization in AWS Organizations. An account can be a member of only one organization at a time.

MES

See [manufacturing execution system](#).

Message Queuing Telemetry Transport (MQTT)

A lightweight, machine-to-machine (M2M) communication protocol, based on the [publish/subscribe](#) pattern, for resource-constrained [IoT](#) devices.

microservice

A small, independent service that communicates over well-defined APIs and is typically owned by small, self-contained teams. For example, an insurance system might include microservices that map to business capabilities, such as sales or marketing, or subdomains, such as purchasing, claims, or analytics. The benefits of microservices include agility, flexible scaling, easy deployment, reusable code, and resilience. For more information, see [Integrating microservices by using AWS serverless services](#).

microservices architecture

An approach to building an application with independent components that run each application process as a microservice. These microservices communicate through a well-defined interface by using lightweight APIs. Each microservice in this architecture can be updated, deployed, and scaled to meet demand for specific functions of an application. For more information, see [Implementing microservices on AWS](#).

Migration Acceleration Program (MAP)

An AWS program that provides consulting support, training, and services to help organizations build a strong operational foundation for moving to the cloud, and to help offset the initial cost of migrations. MAP includes a migration methodology for executing legacy migrations in a methodical way and a set of tools to automate and accelerate common migration scenarios.

migration at scale

The process of moving the majority of the application portfolio to the cloud in waves, with more applications moved at a faster rate in each wave. This phase uses the best practices and lessons learned from the earlier phases to implement a *migration factory* of teams, tools, and processes to streamline the migration of workloads through automation and agile delivery. This is the third phase of the [AWS migration strategy](#).

migration factory

Cross-functional teams that streamline the migration of workloads through automated, agile approaches. Migration factory teams typically include operations, business analysts and owners, migration engineers, developers, and DevOps professionals working in sprints. Between 20 and 50 percent of an enterprise application portfolio consists of repeated patterns that can be optimized by a factory approach. For more information, see the [discussion of migration factories](#) and the [Cloud Migration Factory guide](#) in this content set.

migration metadata

The information about the application and server that is needed to complete the migration. Each migration pattern requires a different set of migration metadata. Examples of migration metadata include the target subnet, security group, and AWS account.

migration pattern

A repeatable migration task that details the migration strategy, the migration destination, and the migration application or service used. Example: Rehost migration to Amazon EC2 with AWS Application Migration Service.

Migration Portfolio Assessment (MPA)

An online tool that provides information for validating the business case for migrating to the AWS Cloud. MPA provides detailed portfolio assessment (server right-sizing, pricing, TCO comparisons, migration cost analysis) as well as migration planning (application data analysis and data collection, application grouping, migration prioritization, and wave planning). The [MPA tool](#) (requires login) is available free of charge to all AWS consultants and APN Partner consultants.

Migration Readiness Assessment (MRA)

The process of gaining insights about an organization's cloud readiness status, identifying strengths and weaknesses, and building an action plan to close identified gaps, using the AWS CAF. For more information, see the [migration readiness guide](#). MRA is the first phase of the [AWS migration strategy](#).

migration strategy

The approach used to migrate a workload to the AWS Cloud. For more information, see the [7 Rs](#) entry in this glossary and see [Mobilize your organization to accelerate large-scale migrations](#).

ML

See [machine learning](#).

modernization

Transforming an outdated (legacy or monolithic) application and its infrastructure into an agile, elastic, and highly available system in the cloud to reduce costs, gain efficiencies, and take advantage of innovations. For more information, see [Strategy for modernizing applications in the AWS Cloud](#).

modernization readiness assessment

An evaluation that helps determine the modernization readiness of an organization's applications; identifies benefits, risks, and dependencies; and determines how well the organization can support the future state of those applications. The outcome of the assessment is a blueprint of the target architecture, a roadmap that details development phases and milestones for the modernization process, and an action plan for addressing identified gaps. For more information, see [Evaluating modernization readiness for applications in the AWS Cloud](#).

monolithic applications (monoliths)

Applications that run as a single service with tightly coupled processes. Monolithic applications have several drawbacks. If one application feature experiences a spike in demand, the entire architecture must be scaled. Adding or improving a monolithic application's features also becomes more complex when the code base grows. To address these issues, you can use a microservices architecture. For more information, see [Decomposing monoliths into microservices](#).

MPA

See [Migration Portfolio Assessment](#).

MQTT

See [Message Queuing Telemetry Transport](#).

multiclass classification

A process that helps generate predictions for multiple classes (predicting one of more than two outcomes). For example, an ML model might ask "Is this product a book, car, or phone?" or "Which product category is most interesting to this customer?"

mutable infrastructure

A model that updates and modifies the existing infrastructure for production workloads. For improved consistency, reliability, and predictability, the AWS Well-Architected Framework recommends the use of [immutable infrastructure](#) as a best practice.

O

OAC

See [origin access control](#).

OAI

See [origin access identity](#).

OCM

See [organizational change management](#).

offline migration

A migration method in which the source workload is taken down during the migration process. This method involves extended downtime and is typically used for small, non-critical workloads.

OI

See [operations integration](#).

OLA

See [operational-level agreement](#).

online migration

A migration method in which the source workload is copied to the target system without being taken offline. Applications that are connected to the workload can continue to function during the migration. This method involves zero to minimal downtime and is typically used for critical production workloads.

OPC-UA

See [Open Process Communications - Unified Architecture](#).

Open Process Communications - Unified Architecture (OPC-UA)

A machine-to-machine (M2M) communication protocol for industrial automation. OPC-UA provides an interoperability standard with data encryption, authentication, and authorization schemes.

operational-level agreement (OLA)

An agreement that clarifies what functional IT groups promise to deliver to each other, to support a service-level agreement (SLA).

operational readiness review (ORR)

A checklist of questions and associated best practices that help you understand, evaluate, prevent, or reduce the scope of incidents and possible failures. For more information, see [Operational Readiness Reviews \(ORR\)](#) in the AWS Well-Architected Framework.

operational technology (OT)

Hardware and software systems that work with the physical environment to control industrial operations, equipment, and infrastructure. In manufacturing, the integration of OT and information technology (IT) systems is a key focus for [Industry 4.0](#) transformations.

operations integration (OI)

The process of modernizing operations in the cloud, which involves readiness planning, automation, and integration. For more information, see the [operations integration guide](#).

organization trail

A trail that's created by AWS CloudTrail that logs all events for all AWS accounts in an organization in AWS Organizations. This trail is created in each AWS account that's part of the organization and tracks the activity in each account. For more information, see [Creating a trail for an organization](#) in the CloudTrail documentation.

organizational change management (OCM)

A framework for managing major, disruptive business transformations from a people, culture, and leadership perspective. OCM helps organizations prepare for, and transition to, new systems and strategies by accelerating change adoption, addressing transitional issues, and driving cultural and organizational changes. In the AWS migration strategy, this framework is called *people acceleration*, because of the speed of change required in cloud adoption projects. For more information, see the [OCM guide](#).

origin access control (OAC)

In CloudFront, an enhanced option for restricting access to secure your Amazon Simple Storage Service (Amazon S3) content. OAC supports all S3 buckets in all AWS Regions, server-side encryption with AWS KMS (SSE-KMS), and dynamic PUT and DELETE requests to the S3 bucket.

origin access identity (OAI)

In CloudFront, an option for restricting access to secure your Amazon S3 content. When you use OAI, CloudFront creates a principal that Amazon S3 can authenticate with. Authenticated principals can access content in an S3 bucket only through a specific CloudFront distribution. See also [OAC](#), which provides more granular and enhanced access control.

ORR

See [operational readiness review](#).

OT

See [operational technology](#).

outbound (egress) VPC

In an AWS multi-account architecture, a VPC that handles network connections that are initiated from within an application. The [AWS Security Reference Architecture](#) recommends setting up your Network account with inbound, outbound, and inspection VPCs to protect the two-way interface between your application and the broader internet.

P

permissions boundary

An IAM management policy that is attached to IAM principals to set the maximum permissions that the user or role can have. For more information, see [Permissions boundaries](#) in the IAM documentation.

personally identifiable information (PII)

Information that, when viewed directly or paired with other related data, can be used to reasonably infer the identity of an individual. Examples of PII include names, addresses, and contact information.

PII

See [personally identifiable information](#).

playbook

A set of predefined steps that capture the work associated with migrations, such as delivering core operations functions in the cloud. A playbook can take the form of scripts, automated runbooks, or a summary of processes or steps required to operate your modernized environment.

PLC

See [programmable logic controller](#).

PLM

See [product lifecycle management](#).

policy

An object that can define permissions (see [identity-based policy](#)), specify access conditions (see [resource-based policy](#)), or define the maximum permissions for all accounts in an organization in AWS Organizations (see [service control policy](#)).

polyglot persistence

Independently choosing a microservice's data storage technology based on data access patterns and other requirements. If your microservices have the same data storage technology, they can encounter implementation challenges or experience poor performance. Microservices are more easily implemented and achieve better performance and scalability if they use the data store best adapted to their requirements.

portfolio assessment

A process of discovering, analyzing, and prioritizing the application portfolio in order to plan the migration. For more information, see [Evaluating migration readiness](#).

predicate

A query condition that returns true or false, commonly located in a WHERE clause.

predicate pushdown

A database query optimization technique that filters the data in the query before transfer. This reduces the amount of data that must be retrieved and processed from the relational database, and it improves query performance.

preventative control

A security control that is designed to prevent an event from occurring. These controls are a first line of defense to help prevent unauthorized access or unwanted changes to your network. For more information, see [Preventative controls](#) in *Implementing security controls on AWS*.

principal

An entity in AWS that can perform actions and access resources. This entity is typically a root user for an AWS account, an IAM role, or a user. For more information, see *Principal* in [Roles terms and concepts](#) in the IAM documentation.

privacy by design

A system engineering approach that takes privacy into account through the whole development process.

private hosted zones

A container that holds information about how you want Amazon Route 53 to respond to DNS queries for a domain and its subdomains within one or more VPCs. For more information, see [Working with private hosted zones](#) in the Route 53 documentation.

proactive control

A [security control](#) designed to prevent the deployment of noncompliant resources. These controls scan resources before they are provisioned. If the resource is not compliant with the control, then it isn't provisioned. For more information, see the [Controls reference guide](#) in the AWS Control Tower documentation and see [Proactive controls](#) in *Implementing security controls on AWS*.

product lifecycle management (PLM)

The management of data and processes for a product throughout its entire lifecycle, from design, development, and launch, through growth and maturity, to decline and removal.

production environment

See [environment](#).

programmable logic controller (PLC)

In manufacturing, a highly reliable, adaptable computer that monitors machines and automates manufacturing processes.

prompt chaining

Using the output of one [LLM](#) prompt as the input for the next prompt to generate better responses. This technique is used to break down a complex task into subtasks, or to iteratively refine or expand a preliminary response. It helps improve the accuracy and relevance of a model's responses and allows for more granular, personalized results.

pseudonymization

The process of replacing personal identifiers in a dataset with placeholder values. Pseudonymization can help protect personal privacy. Pseudonymized data is still considered to be personal data.

publish/subscribe (pub/sub)

A pattern that enables asynchronous communications among microservices to improve scalability and responsiveness. For example, in a microservices-based [MES](#), a microservice can

publish event messages to a channel that other microservices can subscribe to. The system can add new microservices without changing the publishing service.

Q

query plan

A series of steps, like instructions, that are used to access the data in a SQL relational database system.

query plan regression

When a database service optimizer chooses a less optimal plan than it did before a given change to the database environment. This can be caused by changes to statistics, constraints, environment settings, query parameter bindings, and updates to the database engine.

R

RACI matrix

See [responsible, accountable, consulted, informed \(RACI\)](#).

RAG

See [Retrieval Augmented Generation](#).

ransomware

A malicious software that is designed to block access to a computer system or data until a payment is made.

RASCI matrix

See [responsible, accountable, consulted, informed \(RACI\)](#).

RCAC

See [row and column access control](#).

read replica

A copy of a database that's used for read-only purposes. You can route queries to the read replica to reduce the load on your primary database.

re-architect

See [7 Rs](#).

recovery point objective (RPO)

The maximum acceptable amount of time since the last data recovery point. This determines what is considered an acceptable loss of data between the last recovery point and the interruption of service.

recovery time objective (RTO)

The maximum acceptable delay between the interruption of service and restoration of service.

refactor

See [7 Rs](#).

Region

A collection of AWS resources in a geographic area. Each AWS Region is isolated and independent of the others to provide fault tolerance, stability, and resilience. For more information, see [Specify which AWS Regions your account can use](#).

regression

An ML technique that predicts a numeric value. For example, to solve the problem of "What price will this house sell for?" an ML model could use a linear regression model to predict a house's sale price based on known facts about the house (for example, the square footage).

rehost

See [7 Rs](#).

release

In a deployment process, the act of promoting changes to a production environment.

relocate

See [7 Rs](#).

replatform

See [7 Rs](#).

repurchase

See [7 Rs](#).

resiliency

An application's ability to resist or recover from disruptions. [High availability](#) and [disaster recovery](#) are common considerations when planning for resiliency in the AWS Cloud. For more information, see [AWS Cloud Resilience](#).

resource-based policy

A policy attached to a resource, such as an Amazon S3 bucket, an endpoint, or an encryption key. This type of policy specifies which principals are allowed access, supported actions, and any other conditions that must be met.

responsible, accountable, consulted, informed (RACI) matrix

A matrix that defines the roles and responsibilities for all parties involved in migration activities and cloud operations. The matrix name is derived from the responsibility types defined in the matrix: responsible (R), accountable (A), consulted (C), and informed (I). The support (S) type is optional. If you include support, the matrix is called a *RASCI matrix*, and if you exclude it, it's called a *RACI matrix*.

responsive control

A security control that is designed to drive remediation of adverse events or deviations from your security baseline. For more information, see [Responsive controls](#) in *Implementing security controls on AWS*.

retain

See [7 Rs](#).

retire

See [7 Rs](#).

Retrieval Augmented Generation (RAG)

A [generative AI](#) technology in which an [LLM](#) references an authoritative data source that is outside of its training data sources before generating a response. For example, a RAG model might perform a semantic search of an organization's knowledge base or custom data. For more information, see [What is RAG](#).

rotation

The process of periodically updating a [secret](#) to make it more difficult for an attacker to access the credentials.

row and column access control (RCAC)

The use of basic, flexible SQL expressions that have defined access rules. RCAC consists of row permissions and column masks.

RPO

See [recovery point objective](#).

RTO

See [recovery time objective](#).

runbook

A set of manual or automated procedures required to perform a specific task. These are typically built to streamline repetitive operations or procedures with high error rates.

S

SAML 2.0

An open standard that many identity providers (IdPs) use. This feature enables federated single sign-on (SSO), so users can log into the AWS Management Console or call the AWS API operations without you having to create user in IAM for everyone in your organization. For more information about SAML 2.0-based federation, see [About SAML 2.0-based federation](#) in the IAM documentation.

SCADA

See [supervisory control and data acquisition](#).

SCP

See [service control policy](#).

secret

In AWS Secrets Manager, confidential or restricted information, such as a password or user credentials, that you store in encrypted form. It consists of the secret value and its metadata.

The secret value can be binary, a single string, or multiple strings. For more information, see [What's in a Secrets Manager secret?](#) in the Secrets Manager documentation.

security by design

A system engineering approach that takes security into account through the whole development process.

security control

A technical or administrative guardrail that prevents, detects, or reduces the ability of a threat actor to exploit a security vulnerability. There are four primary types of security controls: [preventative](#), [detective](#), [responsive](#), and [proactive](#).

security hardening

The process of reducing the attack surface to make it more resistant to attacks. This can include actions such as removing resources that are no longer needed, implementing the security best practice of granting least privilege, or deactivating unnecessary features in configuration files.

security information and event management (SIEM) system

Tools and services that combine security information management (SIM) and security event management (SEM) systems. A SIEM system collects, monitors, and analyzes data from servers, networks, devices, and other sources to detect threats and security breaches, and to generate alerts.

security response automation

A predefined and programmed action that is designed to automatically respond to or remediate a security event. These automations serve as [detective](#) or [responsive](#) security controls that help you implement AWS security best practices. Examples of automated response actions include modifying a VPC security group, patching an Amazon EC2 instance, or rotating credentials.

server-side encryption

Encryption of data at its destination, by the AWS service that receives it.

service control policy (SCP)

A policy that provides centralized control over permissions for all accounts in an organization in AWS Organizations. SCPs define guardrails or set limits on actions that an administrator can delegate to users or roles. You can use SCPs as allow lists or deny lists, to specify which services

or actions are permitted or prohibited. For more information, see [Service control policies](#) in the AWS Organizations documentation.

service endpoint

The URL of the entry point for an AWS service. You can use the endpoint to connect programmatically to the target service. For more information, see [AWS service endpoints](#) in *AWS General Reference*.

service-level agreement (SLA)

An agreement that clarifies what an IT team promises to deliver to their customers, such as service uptime and performance.

service-level indicator (SLI)

A measurement of a performance aspect of a service, such as its error rate, availability, or throughput.

service-level objective (SLO)

A target metric that represents the health of a service, as measured by a [service-level indicator](#).

shared responsibility model

A model describing the responsibility you share with AWS for cloud security and compliance. AWS is responsible for security *of* the cloud, whereas you are responsible for security *in* the cloud. For more information, see [Shared responsibility model](#).

Shadow AI

Unauthorized [AI](#) applications built or used outside of governed channels within an organization.

SIEM

See [security information and event management system](#).

single point of failure (SPOF)

A failure in a single, critical component of an application that can disrupt the system.

SLA

See [service-level agreement](#).

SLI

See [service-level indicator](#).

SLO

See [service-level objective](#).

split-and-seed model

A pattern for scaling and accelerating modernization projects. As new features and product releases are defined, the core team splits up to create new product teams. This helps scale your organization's capabilities and services, improves developer productivity, and supports rapid innovation. For more information, see [Phased approach to modernizing applications in the AWS Cloud](#).

SPOF

See [single point of failure](#).

star schema

A database organizational structure that uses one large fact table to store transactional or measured data and uses one or more smaller dimensional tables to store data attributes. This structure is designed for use in a [data warehouse](#) or for business intelligence purposes.

strangler fig pattern

An approach to modernizing monolithic systems by incrementally rewriting and replacing system functionality until the legacy system can be decommissioned. This pattern uses the analogy of a fig vine that grows into an established tree and eventually overcomes and replaces its host. The pattern was [introduced by Martin Fowler](#) as a way to manage risk when rewriting monolithic systems. For an example of how to apply this pattern, see [Modernizing legacy Microsoft ASP.NET \(ASMX\) web services incrementally by using containers and Amazon API Gateway](#).

subnet

A range of IP addresses in your VPC. A subnet must reside in a single Availability Zone.

supervisory control and data acquisition (SCADA)

In manufacturing, a system that uses hardware and software to monitor physical assets and production operations.

symmetric encryption

An encryption algorithm that uses the same key to encrypt and decrypt the data.

synthetic testing

Testing a system in a way that simulates user interactions to detect potential issues or to monitor performance. You can use [Amazon CloudWatch Synthetics](#) to create these tests.

system prompt

A technique for providing context, instructions, or guidelines to an [LLM](#) to direct its behavior. System prompts help set context and establish rules for interactions with users.

T

tags

Key-value pairs that act as metadata for organizing your AWS resources. Tags can help you manage, identify, organize, search for, and filter resources. For more information, see [Tagging your AWS resources](#).

target variable

The value that you are trying to predict in supervised ML. This is also referred to as an *outcome variable*. For example, in a manufacturing setting the target variable could be a product defect.

task list

A tool that is used to track progress through a runbook. A task list contains an overview of the runbook and a list of general tasks to be completed. For each general task, it includes the estimated amount of time required, the owner, and the progress.

test environment

See [environment](#).

training

To provide data for your ML model to learn from. The training data must contain the correct answer. The learning algorithm finds patterns in the training data that map the input data attributes to the target (the answer that you want to predict). It outputs an ML model that captures these patterns. You can then use the ML model to make predictions on new data for which you don't know the target.

tool

A function or API that an [agent](#) can invoke to perform operations in external systems.

transit gateway

A network transit hub that you can use to interconnect your VPCs and on-premises networks. For more information, see [What is a transit gateway](#) in the AWS Transit Gateway documentation.

trunk-based workflow

An approach in which developers build and test features locally in a feature branch and then merge those changes into the main branch. The main branch is then built to the development, preproduction, and production environments, sequentially.

trusted access

Granting permissions to a service that you specify to perform tasks in your organization in AWS Organizations and in its accounts on your behalf. The trusted service creates a service-linked role in each account, when that role is needed, to perform management tasks for you. For more information, see [Using AWS Organizations with other AWS services](#) in the AWS Organizations documentation.

tuning

To change aspects of your training process to improve the ML model's accuracy. For example, you can train the ML model by generating a labeling set, adding labels, and then repeating these steps several times under different settings to optimize the model.

two-pizza team

A small DevOps team that you can feed with two pizzas. A two-pizza team size ensures the best possible opportunity for collaboration in software development.

U

uncertainty

A concept that refers to imprecise, incomplete, or unknown information that can undermine the reliability of predictive ML models. There are two types of uncertainty: *Epistemic uncertainty* is caused by limited, incomplete data, whereas *aleatoric uncertainty* is caused by the noise and randomness inherent in the data.

undifferentiated tasks

Also known as *heavy lifting*, work that is necessary to create and operate an application but that doesn't provide direct value to the end user or provide competitive advantage. Examples of undifferentiated tasks include procurement, maintenance, and capacity planning.

upper environments

See [environment](#).

V

vacuuming

A database maintenance operation that involves cleaning up after incremental updates to reclaim storage and improve performance.

version control

Processes and tools that track changes, such as changes to source code in a repository.

VPC peering

A connection between two VPCs that allows you to route traffic by using private IP addresses. For more information, see [What is VPC peering](#) in the Amazon VPC documentation.

vulnerability

A software or hardware flaw that compromises the security of the system.

W

warm cache

A buffer cache that contains current, relevant data that is frequently accessed. The database instance can read from the buffer cache, which is faster than reading from the main memory or disk.

warm data

Data that is infrequently accessed. When querying this kind of data, moderately slow queries are typically acceptable.

window function

A SQL function that performs a calculation on a group of rows that relate in some way to the current record. Window functions are useful for processing tasks, such as calculating a moving average or accessing the value of rows based on the relative position of the current row.

workload

A collection of resources and code that delivers business value, such as a customer-facing application or backend process.

workstream

Functional groups in a migration project that are responsible for a specific set of tasks. Each workstream is independent but supports the other workstreams in the project. For example, the portfolio workstream is responsible for prioritizing applications, wave planning, and collecting migration metadata. The portfolio workstream delivers these assets to the migration workstream, which then migrates the servers and applications.

WORM

See [write once, read many](#).

WQF

See [AWS Workload Qualification Framework](#).

write once, read many (WORM)

A storage model that writes data a single time and prevents the data from being deleted or modified. Authorized users can read the data as many times as needed, but they cannot change it. This data storage infrastructure is considered [immutable](#).

Z

zero-day exploit

An attack, typically malware, that takes advantage of a [zero-day vulnerability](#).

zero-day vulnerability

An unmitigated flaw or vulnerability in a production system. Threat actors can use this type of vulnerability to attack the system. Developers frequently become aware of the vulnerability as a result of the attack.

zero-shot prompting

Providing an [LLM](#) with instructions for performing a task but no examples (*shots*) that can help guide it. The LLM must use its pre-trained knowledge to handle the task. The effectiveness of zero-shot prompting depends on the complexity of the task and the quality of the prompt. See also [few-shot prompting](#).

zombie application

An application that has an average CPU and memory usage below 5 percent. In a migration project, it is common to retire these applications.